



Adaptive machine learning-based surrogate modeling to accelerate PDE-constrained optimization in enhanced oil recovery

Tim Keil¹ · Hendrik Kleikamp¹ · Rolf J. Lorentzen² · Micheal B. Oguntola^{2,3}  · Mario Ohlberger¹

Received: 4 March 2022 / Accepted: 15 September 2022
© The Author(s) 2022

Abstract

In this contribution, we develop an efficient surrogate modeling framework for simulation-based optimization of enhanced oil recovery, where we particularly focus on polymer flooding. The computational approach is based on an adaptive training procedure of a neural network that directly approximates an input-output map of the underlying PDE-constrained optimization problem. The training process thereby focuses on the construction of an accurate surrogate model solely related to the optimization path of an outer iterative optimization loop. True evaluations of the objective function are used to finally obtain certified results. Numerical experiments are given to evaluate the accuracy and efficiency of the approach for a heterogeneous five-spot benchmark problem.

Keywords PDE-constrained optimization · Enhanced oil recovery · Machine learning · Neural networks · Surrogate modeling · Ensemble-based optimization

Mathematics Subject Classification (2010) 49M41 · 68T07 · 90C90

1 Introduction

Water flooding remains the most frequently used secondary oil recovery method. However, the percentage of original oil in place left after the cessation of water flooding in many reservoir fields is estimated to be as high as 50–70% [1–3]. The reduced performance of water flooding leading to the sizable leftover of oil has been linked to many factors such as the presence of unfavorable mobility ratios (due to

Communicated by: Gianluigi Rozza

✉ Micheal B. Oguntola
micheal.b.oguntola@uis.no

Extended author information available on the last page of the article

heavy oil) and high level of heterogeneity (in porosity and permeability) in the reservoir [4]. For these reasons, enhanced oil recovery (EOR) methods are employed to improve the performance of water flooding in order to increase oil production and minimize environmental stress.

Polymer flooding is a matured chemical EOR method, suitable for heavy oil reservoir development, with over four decades of practical applications [5, 6]. It involves injecting long chains of high-molecular-weight soluble polymers along with water flooding. The polymer EOR mechanism includes reducing mobility ratios of the oil-water system and early water breakthrough in the reservoir by increasing the viscosity of injected water and consequently improving vertical and aerial sweep efficiencies of the injected fluid.

The EOR process of polymer flooding can significantly increase the oil production [6]. However, compared to water flooding, the operational cost and the risk associated with polymer flooding are higher. More so, since injecting more than necessary polymer into the reservoir can lead to insignificant oil increment, it is imperative to optimize the injection strategy of polymer flooding for field application to avoid unnecessarily high operational costs with no profit.

Conventionally, a reservoir simulation model is combined with a numerical optimization technique to determine an optimal control (including water rates, polymer concentrations of injection wells, liquid rates, or bottom hole pressures of production wells) for polymer flooding. The aim is to maximize a given reservoir performance measure (RPM), such as the total oil production or the net present value (NPV) function over the reservoir life. The simulation model is usually a complex numerical reservoir simulator that requires substantial data accounting for geology and geometry of the reservoir or rock and fluid properties. In this study, the model simulates the oil reservoir response (in the form of fluid production) to a given polymer flooding control per time. On this account, we estimate the RPM of a given control strategy.

Further, the complexity of a reservoir simulator leads to a high computational effort for simulating a given polymer flooding scenario. It contributes to the inefficiency of gradient-based solution techniques (e.g., the ensemble-based optimization (EnOpt) method) for polymer EOR optimization problems, since the (approximate) gradient of the objective function with respect to the control variables requires several function evaluations, with each relying on a time-consuming polymer model simulation [7–9]. More so, for large-scale polymer problems discretized into a large number of grid cells, a single model evaluation may take several hours to complete. For this reason, we propose a machine learning-based approach to approximate the computationally demanding objective function.

In classical approaches of model order reduction or surrogate modeling, the expensive evaluation of the objective function due to the PDE constraints is replaced by an a priori trained surrogate that can be efficiently evaluated with respect to the optimization parameters, e.g., by reduced basis methods [10–14]. These works share the fact that the offline time for constructing the respective surrogate can be ignored since the respective optimization parameters do not necessarily belong to the reduction parameters. If instead a single optimization problem is to be solved and no pre-trained surrogate model exists (as is the case in our application) the offline time

for constructing a globally accurate surrogate can be large or, in the case of a high dimensional parameter space, even prohibitive. Furthermore, reliable error estimation is required to ensure a desired accuracy of the surrogate.

In this work, we make use of an adaptive surrogate modeling approach, where a surrogate model is constructed during the outer optimization loop through adaptive learning that is targeted towards an accurate input-output map in the vicinity of the chosen parameters during the optimization loop. The overall algorithm thus combines costly full order model (FOM) evaluations, training of machine learning (ML) based surrogate models, as well as evaluations of the successively trained ML models. In model reduction for parameterized systems [15], such adaptive enrichment approaches have been recently proposed and successfully applied in the context of PDE-constrained parameter optimization, e.g., by a progressive construction of reduced-order models (ROMs) [16, 17], or in combination with trust-region optimization [18–20]. Recently, first ideas were presented to combine online enrichment for ROMs with machine learning-based surrogate modeling [21–23]. In this contribution, we use feedforward deep neural networks (DNNs) to obtain surrogate models of the underlying input-output map that directly map the optimization parameters to the output of the objective function. For instance in [24, 25], DNNs have also been applied as ROM for parametrized PDEs in combination with reduced basis methods. Deep convolutional autoencoders, a special type of neural networks, were used in [26] to obtain a nonlinear approximation manifold that replaces the linear subspace from the reduced basis method.

Artificial neural networks also gained attention in the context of enhanced oil recovery in recent years, see [27–29], for instance. However, these approaches mainly focus on accelerating the evaluation of the costly objective function without providing a way to solve polymer EOR optimization problems using the proposed surrogate models. In [30], the authors describe an algorithm to obtain a global surrogate model that is applied as a replacement for the objective function in a genetic algorithm. The global approximation of the objective is computed a priori before applying the optimization routine. In [31], artificial neural networks are employed to facilitate the decision process for a specific EOR method.

Concerning acceleration of PDE-constrained optimization in general, DNNs are, for instance, used in [32] to replace costly simulations within the optimization loops by evaluations of surrogate models. The main idea of the *ISMO* algorithm described in [32] is to run multiple parallel optimization routines starting from different initial guesses and to construct DNN surrogate models using training data collected at the final iterates of these optimization algorithms. The training data is computed by costly evaluations of the exact objective function (involving the solution of PDEs). In contrast, the optimization routines use the respective surrogate model to speed up the computations. Iteratively, a surrogate model is built to approximate the true objective function near local optima. The approximation quality also serves as the stopping criterion of the algorithm. Another approach involving physics-informed deep operator networks to accelerate PDE-constrained optimization in a self-supervised manner has recently been suggested in [33].

The idea of not having a global surrogate model, but only approximations of the objective function that are locally accurate, is also one of the main motivations for

our algorithm. In contrast to the procedure in [32] described previously, we iteratively construct DNN surrogate models tailored towards the objective function along a single optimization path. We consider only a single initial guess but check for convergence by taking into account the true objective function. This stopping criterion certifies that the resulting control is approximately a (local) optimum of the true objective function and not only of the surrogate. Further, we do not assume that the derivative of the DNN surrogate with respect to its inputs is available but reuse the EnOpt procedure when optimizing with the surrogate model.

The remainder of this article is organized as follows. In Section 2 we introduced the polymer flooding model for EOR and formulate an optimization problem for the economic value of the reservoir response. Section 3 introduces a classical ensemble-based optimization algorithm based on a FOM approximation of the polymer flooding model. Feedforward DNNs to approximate the input-output map are introduced in Section 4. In Section 5, we finally present and discuss our new adaptive FOM-ML-based optimization algorithm, which is evaluated numerically for a five-spot benchmark problem in Section 6. Last but not least, a conclusion and outlook are given in Section 7.

2 Optimization of polymer flooding in enhanced oil recovery

The problem of predicting the optimal injection strategy of the polymer EOR method can be formulated as a constrained optimization problem. The setup involves solving a maximization problem in which the objective function, the RPM, is defined on a given set of controllable variables. For the polymer EOR method, a complete set of control variables includes the concentration (and hence volume size of the polymer) and control variables (such as water injection rate, oil production rate, and/or bottom hole pressure for the injecting or producing wells) for water flooding over the producing lifespan of the reservoir.

2.1 Polymer flooding model

As mentioned in the introduction, the optimization process is usually performed on a simulation model of the real reservoir [34]. Here, we consider a polymer flooding simulation model, which is an extension of the black-oil model with a continuity equation for the polymer component [9, 35]. The black-oil model is a special multi-component multi-phase flow model with no diffusion among the fluid components [36]. It assumes that all hydrocarbon species are considered as two components, namely, oil and gas at surface conditions, and can be partially or entirely dissolved in each other to form the oil and gas phases. Further, there is an aqueous phase that consists of only one component called water.

For brevity, we first state the polymer flooding model without mentioning its dependence on the controls and geological parameters explicitly. Hence, in what follows, we assume that fixed sets of controls and geological parameters are given.

In the polymer model, usually, it is assumed that polymer forms an additional component transported in the aqueous phase of the black-oil model and has no effect on the oil phase. We identify those quantities associated with the water, oil, gas, and polymer components with subscripts “W”, “O”, “G”, and “P”. In general, the polymer model consists of the following system of partial differential equations:

$$\frac{\partial}{\partial t}(\phi b_W s_W) + \nabla \cdot b_W \mathbf{v}_W = q_W, \tag{1a}$$

$$\frac{\partial}{\partial t} \phi (b_O s_O + r_{OG} b_G s_G) + \nabla \cdot (b_O \mathbf{v}_O + r_{OG} b_G \mathbf{v}_G) = q_O, \tag{1b}$$

$$\frac{\partial}{\partial t} \phi (b_G s_G + r_{GO} b_O s_O) + \nabla \cdot (b_G \mathbf{v}_G + r_{GO} b_O \mathbf{v}_O) = q_G, \tag{1c}$$

$$\frac{\partial}{\partial t} \left[\phi (1 - s_{ipv}) s_W + \frac{\rho_r c_a}{b_W c} (1 - \phi) \right] + \nabla \cdot \mathbf{v}_P = q_W, \tag{1d}$$

where ϕ is the rock porosity, s_α , b_α , q_α , and \mathbf{v}_α denote the (unknown) saturation, inverse formation-volume factor (depending on the respective density ρ_α), volumetric source (flow rate per unit volume), and Darcy’s flux of phase $\alpha \in \{W, O, G\}$, and r_{OG} and r_{GO} denote the oil-gas and gas-oil ratios. The quantities $\mathbf{v}_P, c_a, s_{ipv}$, and c denote the Darcy’s flux, adsorption concentration, inaccessible pore volume, and concentration of the polymer solution, and ρ_r is the density of the reservoir rock.

In addition to the system (1), empirical closure equations for relative permeabilities and capillary pressure in three-phase flow in porous media are applied. Here, the unknown primary variables are phase saturations s_α (or component accumulations) and pressures p_α , and thus, appropriate initial and boundary conditions are defined.

Based on the type of injection and/or production well (e.g., vertical, horizontal, or multi-segment), a suitable well model [37, 38] is coupled with (1) to measure the volumetric flow rates, which depend on the state of the reservoir. A standard well model for vertical wells is given as follows.

The volumetric flow rates q_α for $\alpha \in \{W, O, G\}$ in a multi-phase polymer model are computed using a semi-analytical model according to [38, 39] and are given by

$$q_W = \frac{k_{RW}(s_W)}{\mu_{W,eff} R_k(c)} WI(p_{bh} - p_W - \rho_W g(z_{bh} - z)), \tag{2a}$$

$$q_O = \frac{k_{RO}(s_O)}{\mu_{O,eff}} WI(p_{bh} - p_O - \rho_O g(z_{bh} - z)), \tag{2b}$$

$$q_G = \frac{k_{RG}(s_G)}{\mu_{G,eff}} WI(p_{bh} - p_G - \rho_G g(z_{bh} - z)). \tag{2c}$$

Here, $k_{R\alpha}(s_\alpha)$, ρ_α , p_α , and $\mu_{\alpha,eff}$ are the saturation-dependent relative permeability, density, pressure, and effective viscosity of phase $\alpha \in \{W, O, G\}$, WI is the well index, z_{bh} is the well datum level depth, p_{bh} is the bottom hole pressure at the well datum level, z is the depth, $R_k(c)$ models the reduced permeability experienced by the water-polymer mixture, and g is the magnitude of the gravitational acceleration.

Individual wells are usually controlled by surface flow rates or bottom hole pressures. Additional equations which enforce limit values for the component rates and bottom hole pressures are

$$\begin{aligned} p_{bh} - p_{bh}^{limit} &\leq 0, \\ q_\alpha - q_\alpha^{limit} &\leq 0, \end{aligned}$$

where q_α^{limit} is the desired surface-volume rate limit for component α , e.g., field oil rate at the production well, and p_{bh}^{limit} is the desired bottom hole pressure limit. Also, logic constraints to determine what happens if the computed rates or pressures violate the operational constraints, in which case a well may switch from rate control to pressure control, etc., are imposed.

If $q_{\alpha,i}$ is the field volumetric flow rate (in sm^3/day) of component $\alpha \in \{W,O,G\}$ in the production wells over the time interval Δt_i , the field production total (in sm^3) of the component α is given as $Q_{\alpha P,i} = q_{\alpha,i} \Delta t_i$. For polymer production total (in kg), $Q_{PP,i} = c_L q_{W,i} \Delta t_i$, where c_L is the leftover field polymer concentration (in kg/sm^3) after adsorption. Injection quantities $Q_{PI,i}$ and $Q_{WI,i}$ are computed similarly, however with volumetric flow rates in the injection wells.

As already mentioned above, the solution of the polymer flooding model stated in (1) depends on a given control vector \mathbf{u} , see Section 2.2 for a detailed description of the components of the control vector, and a set of geological properties θ . Consequently, all involved unknowns depend on \mathbf{u} and θ and the same holds for q_w , q_o , and q_g . From now on, we thus write $Q_{\alpha P,i}(\mathbf{u}, \theta)$ for the field production total of component $\alpha \in \{W,O,G\}$, depending on the controls \mathbf{u} and the parameters θ , within the time interval Δt_i , similar as above. We further write $Q_{PP}(\mathbf{u}, \theta)$ for the polymer production total, and $Q_{PI,i}(\mathbf{u}, \theta)$ and $Q_{WI,i}(\mathbf{u}, \theta)$ for the polymer and water injection.

2.2 Optimization of the economic value of the reservoir response

This study considers the annually discounted net present value (NPV) function as the RPM, similar to the one in [7, 40]. The NPV function is related to the control variables through the polymer simulation model (1). For every polymer control strategy, the NPV function evaluates the economic value of the reservoir response. Also, because the injection and production facilities have limited capacity, the control variables are subject to bound constraints.

Suppose that the geological properties of the oil reservoir of interest, such as porosity and permeability are known and denoted by θ . Let $\mathcal{D} = \mathbb{R}^{N_u}$ be the domain of control vectors of polymer flooding for the reservoir, such that

$$\mathbf{u} = \left[u_1^1, u_2^1, \dots, u_{N_w}^1, \dots, u_1^{N_t}, u_2^{N_t}, \dots, u_{N_w}^{N_t} \right]^T,$$

where T means transpose. The subscript of each component of \mathbf{u} denotes the well index, the superscript is the control time step, N_w and N_t denote the number of wells and time steps for each well, respectively, and $N_u = N_w \cdot N_t$ is the total number of control variables. Each component u_j^i in \mathbf{u} represents a control type (e.g., polymer

concentration or injection rate, oil or water rate, bottom hole pressure) of well j at the time step i .

The N_u -dimensional optimization problem for polymer flooding is to find the optimal $\mathbf{u} \in \mathcal{D}$ that maximizes the NPV function subject to bound constraints. That is

$$\underset{\mathbf{u} \in \mathcal{D}}{\text{maximize}} \quad J(\mathbf{u}, \boldsymbol{\theta}) := \sum_{i=1}^{N_t} \frac{J_i(\mathbf{u}, \boldsymbol{\theta})}{(1+d_\tau)^{\frac{t_i}{\tau}}} \tag{3a}$$

with

$$\begin{aligned} J_i(\mathbf{u}, \boldsymbol{\theta}) &:= r_{OP}Q_{OP,i}(\mathbf{u}, \boldsymbol{\theta}) + r_{GP}Q_{GP,i}(\mathbf{u}, \boldsymbol{\theta}) - R_i(\mathbf{u}, \boldsymbol{\theta}), \\ R_i(\mathbf{u}, \boldsymbol{\theta}) &:= r_{WI}Q_{WI,i}(\mathbf{u}, \boldsymbol{\theta}) + r_{WP}Q_{WP,i}(\mathbf{u}, \boldsymbol{\theta}) \\ &\quad + r_{PI}Q_{PI,i}(\mathbf{u}, \boldsymbol{\theta}) + r_{PP}Q_{PP,i}(\mathbf{u}, \boldsymbol{\theta}), \end{aligned} \tag{3b}$$

subject to

$$u_j^{\text{low}} \leq u_j^i \leq u_j^{\text{upp}} \quad \text{for all } j = 1, \dots, N_w, \quad i = 1, \dots, N_t, \tag{3c}$$

where J_i denotes the cumulative NPV value in the i -th simulation time step. Further, d_τ is the discount rate for a period of τ days, t_i is the cumulative time (in days) starting from the beginning of production up to the i -th time step, and $\Delta t_i := t_i - t_{i-1}$ is the time difference (in days) between the time steps t_i and t_{i-1} . The scalars r_{OP}, r_{GP}, r_{WI} and r_{WP} denote the prices of oil and gas production and the cost of handling water injection and production (in USD/sm³) respectively, and r_{PI} and r_{PP} are the costs of polymer injection and production (in USD/kg). In addition, $Q_{WI,i}$ and $Q_{PI,i}$ are the total water injection (in sm³) and total polymer injection or slug size (in kg) over the time interval Δt_i . The quantities $Q_{OP,i}$, $Q_{WP,i}$ and $Q_{GP,i}$ denote the total oil, water and gas productions (in sm³) over the time interval Δt_i , while $Q_{PP,i}$ represents the total polymer production (in kg) over the time interval Δt_i . The quantities $Q_{OP,i}$, $Q_{WI,i}$, $Q_{WP,i}$, $Q_{GP,i}$, $Q_{PI,i}$ and $Q_{PP,i}$ are computed at each control time step i for given \mathbf{u} and fixed $\boldsymbol{\theta}$ from the polymer flooding model (1) and the well equations (2).

The evaluation of the objective function J in (3b) shall be referred to as the full order model (FOM) function evaluation in the remainder of this study. Therefore, the constrained optimization problem presented in (3) can be interpreted as the FOM optimization problem for polymer flooding, given a suitable discretization of the system (1) (see Section 6.1 for details). Also, because $\boldsymbol{\theta}$ is fixed during the optimization process, J is considered a function of \mathbf{u} only, and hence we often write $J(\mathbf{u})$ and $J_i(\mathbf{u})$. The solution method utilized for this optimization problem is presented in the next section.

3 Ensemble-based optimization algorithm

In this work, the FOM solution to problem (3) follows from the application of the adaptive ensemble-based optimization (EnOpt) method analogous to the one presented in [7, 41, 42]. We again emphasize that we restrict our attention to a fixed

choice of geological parameters θ . Since we apply the EnOpt algorithm later on in our surrogate-based algorithm to a function different from J , we subsequently begin by describing the algorithm in its general form. Afterwards, we discuss the application of the EnOpt algorithm to the objective function J and the resulting computational costs.

3.1 Optimization algorithm for a general objective function

In what follows, we describe the EnOpt algorithm for a general objective function $F : \mathbb{R}^{N_u} \rightarrow \mathbb{R}$ to iteratively solve the optimization problem

$$\underset{\mathbf{u} \in \mathcal{D}}{\text{maximize}} \quad F(\mathbf{u}) \quad (4a)$$

$$\text{subject to} \quad u_j^{\text{low}} \leq u_j^i \leq u_j^{\text{upp}} \quad \text{for all } j = 1, \dots, N_w, \quad i = 1, \dots, N_t. \quad (4b)$$

The EnOpt method is an iterative method in which one starts with an initial guess \mathbf{u}_0 that is usually based on experimental facts in such a way that the underlying constraints in (4b) are satisfied. We sequentially seek for an improved approximate solution \mathbf{u} that maximizes $F(\mathbf{u})$ using a preconditioned (with covariance matrix adaptation) gradient ascent method given by

$$\hat{\mathbf{u}}_{k+1} = \mathbf{u}_k + \beta_k \mathbf{d}_k, \quad (5)$$

$$\mathbf{d}_k \approx \frac{\mathbf{C}_{\mathbf{u}_k}^k \mathbf{G}_k}{\|\mathbf{C}_{\mathbf{u}_k}^k \mathbf{G}_k\|_\infty}, \quad (6)$$

where $k = 0, 1, 2, \dots$ denotes the optimization iteration. The tuning parameter β_k for the step size is computed using an auxiliary line search [43] and is selected such that $0 < \beta_k \leq 1$. Furthermore, $\mathbf{C}_{\mathbf{u}_k}^k \in \mathbb{R}^{N_u \times N_u}$ denotes the user-defined covariance matrix of the control variables at the k -th iteration and $\mathbf{G}_k \in \mathbb{R}^{N_u}$ is the approximate gradient of F with respect to the control variables, preconditioned with $\mathbf{C}_{\mathbf{u}_k}^k$ to obtain the search direction at iteration k .

To ensure that the constraints in (5) are satisfied, the original solution domain of the control variables is projected to the set of admissible controls \mathcal{D}_{ad} , defined as

$$\mathcal{D}_{\text{ad}} := \{\mathbf{u} \in \mathcal{D} : u_j^{\text{low}} \leq u_j^i \leq u_j^{\text{upp}} \text{ for all } j = 1, \dots, N_w, \quad i = 1, \dots, N_t\}, \quad (7)$$

which corresponds to the constraints in (4b). The updating scheme in (5) is performed in \mathcal{D}_{ad} . We utilize a component-wise projection $P_{\mathcal{D}_{\text{ad}}} : \mathcal{D} \rightarrow \mathcal{D}_{\text{ad}}$ on the update $\hat{\mathbf{u}}_{k+1} \in \mathcal{D}$, such that

$$\mathbf{u}_{k+1} = P_{\mathcal{D}_{\text{ad}}}(\hat{\mathbf{u}}_{k+1}) \in \mathcal{D}_{\text{ad}}. \quad (8)$$

In practical applications, it is not common to have controls at different wells to correlate, but the controls may vary smoothly with time at individual wells.

Hence, the use of $\mathbf{C}_{\mathbf{u}_k}^k$ in Equation (5) enforces this regularization on the control updates. At $k = 0$, we utilize a temporal covariance function given by

$$\text{Cov}\left(u_j^i, u_j^{i+h}\right) = \sigma_j^2 \rho^h \left(\frac{1}{1-\rho^2}\right), \quad \text{for all } h \in \{0, \dots, N_t - i\}, \tag{9}$$

from a stationary auto regression of order 1 (i.e., AR(1)) model [44] to compute $\mathbf{C}_{\mathbf{u}_0}^0$ with an assumption that controls of different wells are uncorrelated. The variance for the well j is given by $\sigma_j^2 > 0$, and $\rho \in (-1, 1)$ is the correlation coefficient used to introduce a level of dependence between controls of individual wells at different control time steps (since the AR(1) model is stationary).

The formulation above gives rise to a block diagonal matrix $\mathbf{C}_{\mathbf{u}_0}^0$, which is updated by matrices with rank one at subsequent iterations, using the statistical method presented in [42], to obtain an improved covariance matrix $\mathbf{C}_{\mathbf{u}_k}^k$. For this reason, the solution method in Equation (5) is referred to as the adaptive EnOpt algorithm.

We compute the preconditioned approximate gradient $\mathbf{C}_{\mathbf{u}_k}^k \mathbf{G}_k$ following the approach of the standard EnOpt algorithm. At the k -th iteration, we sample $N \in \mathbb{N}$ control vectors $\mathbf{u}_{k,m} \in \mathcal{D}_{\text{ad}}$, for $m = 1, \dots, N$, from a multivariate Gaussian distribution with mean equal to the k -th control vector \mathbf{u}_k and covariance matrix given by $\mathbf{C}_{\mathbf{u}_k}^k$. Here, the additional subscript m is used to differentiate the perturbed control vectors from the one obtained by Equation (5). The cross-covariance of the control vector \mathbf{u}_k and the objective function $F(\mathbf{u}_k)$ at the k -th iteration is approximated according to [45] as

$$\mathbf{C}_{\mathbf{u}_k, F}^k := \frac{1}{N-1} \sum_{m=1}^N (\mathbf{u}_{k,m} - \mathbf{u}_k)(F(\mathbf{u}_{k,m}) - F(\mathbf{u}_k)). \tag{10}$$

Since $\mathbf{u}_{k,m} \sim \mathcal{N}(\mathbf{u}_k, \mathbf{C}_{\mathbf{u}_k}^k)$ for $m = 1, \dots, N$, we assume in Equation (10) that the mean of $\{\mathbf{u}_{k,m}\}_{m=1}^N$ is approximated by \mathbf{u}_k . By first-order Taylor series expansion of F about \mathbf{u}_k , it can easily be deduced that Equation (10) is an approximation of $\mathbf{C}_{\mathbf{u}_k}^k \mathbf{G}_k$ at the k -th iteration, that is

$$\mathbf{C}_{\mathbf{u}_k}^k \mathbf{G}_k \approx \mathbf{C}_{\mathbf{u}_k, F}^k, \tag{11}$$

see [41, 46] for a detailed proof. Therefore, we choose the search direction as $\mathbf{d}_k = \mathbf{C}_{\mathbf{u}_k, F}^k / \|\mathbf{C}_{\mathbf{u}_k, F}^k\|_\infty$ in Equation (5). The updating scheme in Equation (5) is performed until the convergence criterion

$$F(\mathbf{u}_k) \leq F(\mathbf{u}_{k-1}) + \epsilon \tag{12}$$

is satisfied, where $\epsilon > 0$ is a specified tolerance.

For an arbitrary objective function F , the EnOpt procedure is summarized in Algorithm 1. In this algorithm, the OPTSTEP function replicates a single optimization step in the EnOpt procedure and is detailed in Algorithm 2. We note that

returning the set of function values T_{k+1} does not play a role in Algorithm 1 but is crucial for training the surrogate model in Section 5. The line search procedure `LINESEARCH` can be found in Algorithm 3.

Algorithm 1 EnOpt algorithm

Require: function $F: \mathbb{R}^{N_u} \rightarrow \mathbb{R}$ for which to solve (4); initial guess $\mathbf{u}_0 \in \mathbb{R}^{N_u}$, sample size $N \in \mathbb{N}$, tolerance $\varepsilon > 0$, maximum number of iterations k^* , initial step size $\beta > 0$, step size contraction $r \in (0, 1)$, maximum number of step size trials $v^* \in \mathbb{N}$

Ensure: approximate solution $\mathbf{u}^* \in \mathbb{R}^{N_u}$ of (4)

- 1: **function** ENOPT[F]($\mathbf{u}_0, N, \varepsilon, k^*, \beta, r, v^*$)
- 2: $\mathbf{u}_1, T_1 \leftarrow \text{OPTSTEP}[F](\mathbf{u}_0, N, 0, \beta, r, v^*)$
- 3: $k \leftarrow 1$
- 4: **while** $F(\mathbf{u}_k) > F(\mathbf{u}_{k-1}) + \varepsilon$ and $k < k^*$ **do**
- 5: $\mathbf{u}_{k+1}, T_{k+1} \leftarrow \text{OPTSTEP}[F](\mathbf{u}_k, N, k, \beta, r, v^*)$
- 6: $k \leftarrow k + 1$
- 7: **end while**
- 8: **return** $\mathbf{u}^* \leftarrow \mathbf{u}_k$
- 9: **end function**

Algorithm 2 OptStep algorithm

Require: function $F: \mathbb{R}^{N_u} \rightarrow \mathbb{R}$; current control vector $\mathbf{u}_k \in \mathbb{R}^{N_u}$, sample size $N \in \mathbb{N}$, number of iteration k , initial step size $\beta > 0$, step size contraction $r \in (0, 1)$, maximum number of step size trials $v^* \in \mathbb{N}$

Ensure: update $\mathbf{u}_{k+1} \in \mathbb{R}^{N_u}$ of the controls, set T_{k+1} of N pairs of the form $(\mathbf{u}, F(\mathbf{u}))$

- 1: **function** OPTSTEP[F]($\mathbf{u}_k, N, k, \beta, r, v^*$)
- 2: **if** $k = 0$ **then**
- 3: Compute the initial covariance matrix $\mathbf{C}_{\mathbf{u}_0}^0$ using (9)
- 4: **else**
- 5: Compute the covariance matrix $\mathbf{C}_{\mathbf{u}_k}^k$ using the formulation in [43]
- 6: **end if**
- 7: Sample N control vectors $\{\mathbf{u}_{k,j}\}_{j=1}^N$ from a distribution $\mathcal{N}(\mathbf{u}_k, \mathbf{C}_{\mathbf{u}_k}^k)$
- 8: Compute vector $\mathbf{C}_{\mathbf{u}_k, F}^k$ according to (10) and store values $\{F(\mathbf{u}_{k,j})\}_{j=1}^N$
- 9: Compute the search direction $\mathbf{d}_k = \mathbf{C}_{\mathbf{u}_k, F}^k / \|\mathbf{C}_{\mathbf{u}_k, F}^k\|_\infty$
- 10: $\mathbf{u}_{k+1} \leftarrow \text{LINESEARCH}[F](\mathbf{u}_k, \mathbf{d}_k, \beta, r, v^*)$
- 11: $T_{k+1} \leftarrow \{(\mathbf{u}_{k,j}, F(\mathbf{u}_{k,j}))\}_{j=1}^N$
- 12: **return** $\mathbf{u}_{k+1}, T_{k+1}$
- 13: **end function**

Algorithm 3 Line search

Require: function $F: \mathbb{R}^{N_u} \rightarrow \mathbb{R}$; current controls $\mathbf{u}_k \in \mathbb{R}^{N_u}$, search direction $\mathbf{d}_k \in \mathbb{R}^{N_u}$, initial step size $\beta > 0$, step size contraction $r \in (0, 1)$, maximum number of step size trials $\nu^* \in \mathbb{N}$, tolerance $\varepsilon > 0$

Ensure: update $\mathbf{u}_{k+1} \in \mathbb{R}^{N_u}$ of the controls

- 1: **function** LINESEARCH[F]($\mathbf{u}_k, \mathbf{d}_k, \beta, r, \nu^*$)
- 2: $\beta_k \leftarrow \beta$
- 3: Compute \mathbf{u}_{k+1} according to (8)
- 4: $\nu \leftarrow 0$
- 5: **while** $F(\mathbf{u}_{k+1}) - F(\mathbf{u}_k) \leq \varepsilon$ and $\nu < \nu^*$ **do**
- 6: $\beta_k \leftarrow r \beta_k$
- 7: Compute \mathbf{u}_{k+1} according to (8)
- 8: $\nu \leftarrow \nu + 1$
- 9: **end while**
- 10: **return** \mathbf{u}_{k+1}
- 11: **end function**

3.2 FOM-EnOpt algorithm for enhanced oil recovery

Eventually, we are interested in solving the optimization problem (3) for polymer flooding in enhanced oil recovery. As already discussed in the introduction, our contribution is concerned with the development of a surrogate-based algorithm to reduce the computational costs for solving (3). To this end, if the EnOpt algorithm is used to maximize the function J , defined in Equation (3b), we refer to Algorithm 1 as the *FOM-EnOpt algorithm*. That is, the FOM-EnOpt algorithm is given as ENOPT[J], see Algorithm 4.

Algorithm 4 FOM-EnOpt algorithm

Require: initial guess $\mathbf{u}_0 \in \mathbb{R}^{N_u}$, sample size $N \in \mathbb{N}$, tolerance $\varepsilon > 0$, maximum number of iterations $k^* \in \mathbb{N}$, initial step size $\beta > 0$, step size contraction $r \in (0, 1)$, maximum number of step size trials $\nu^* \in \mathbb{N}$

Ensure: approximate solution $\mathbf{u}^* \in \mathbb{R}^{N_u}$ of (3)

- 1: **function** FOM-ENOPT($\mathbf{u}_0, N, \varepsilon, k^*, \beta, r, \nu^*$)
- 2: **return** ENOPT[J]($\mathbf{u}_0, N, \varepsilon, k^*, \beta, r, \nu^*$)
- 3: **end function**

As already indicated, we are concerned with the computational effort of the FOM-EnOpt algorithm. Let us recall that evaluating J as in (3b) has the complexity of the high-fidelity reservoir simulator, which, in itself, requires the solution of the discretized polymer flooding model equations (1). In Algorithm 1, the most expensive part is to call OPTSTEP[J], which requires N evaluations of J in Line 8 of Algorithm 2 such that the direction \mathbf{d}_k can be computed in Line 9. Furthermore, the line search in Line 10 evaluates J for every search step. Suppose the simulation time for computing J is particularly large. In that case, the

FOM-EnOpt algorithm can be extremely costly, especially if many optimization steps are required since $\text{OPTSTEP}[J]$ is called at every iteration step. In this case, all steps in Algorithm 1 and Algorithm 2 that do not require evaluating J are computationally negligible.

Since expensive FOM evaluations are very likely to happen for the presented application, we aim to derive a surrogate-based algorithm that uses an approximation of J whenever possible and thus tries to reduce the number of calls of $\text{OPTSTEP}[J]$. Instead, FOM information is reused whenever possible and only computed when necessary. The following section introduces a machine learning-based way for deriving suitable non-intrusive surrogate models.

4 Neural networks as surrogate model for the input-output map

Deep neural networks (DNNs) are machine learning algorithms suitable for approximating functions without knowing their exact structure. Instead, DNNs can be fitted to approximately reproduce known target values for a set of given inputs, and thus, learn from examples of labeled data.

A particular class of DNNs are *feedforward neural networks*, in which no cyclic flow of information is allowed. This study considers feedforward neural networks consisting of (fully connected) linear layers combined with a nonlinear activation function. Our description of these types of DNNs is based on formal definitions that can be found in [47] and [48], for instance.

Feedforward neural networks are used to approximate a given function $f: \mathbb{R}^{N_{\text{in}}} \rightarrow \mathbb{R}^{N_{\text{out}}}$ for a certain input dimension $N_{\text{in}} \in \mathbb{N}$ and an output dimension $N_{\text{out}} \in \mathbb{N}$. To this end, let $L \in \mathbb{N}$ denote the *number of layers* in the neural network, and $N_{\text{in}} = N_0, N_1, \dots, N_{L-1}, N_L = N_{\text{out}} \in \mathbb{N}$ the *numbers of neurons* in each layer. Furthermore, the *weights* and *biases* in layer $i \in \{1, \dots, L\}$ are denoted by $W_i \in \mathbb{R}^{N_i \times N_{i-1}}$ and $b_i \in \mathbb{R}^{N_i}$. We assemble the weights and biases in an L -tuple $\mathbf{W} = ((W_1, b_1), \dots, (W_L, b_L))$. Moreover, let $\rho: \mathbb{R} \rightarrow \mathbb{R}$ be the so-called *activation function* and $\rho_n^*: \mathbb{R}^n \rightarrow \mathbb{R}^n$ the component-wise application of the activation function ρ for dimension $n \in \mathbb{N}$, that is $\rho_n^*(y) := [\rho(y_1), \dots, \rho(y_n)]^T \in \mathbb{R}^n$ for $y \in \mathbb{R}^n$. Then we can define the corresponding feedforward neural network in the following way: The *feedforward neural network* with weights and biases \mathbf{W} and activation function ρ for approximating $f: \mathbb{R}^{N_{\text{in}}} \rightarrow \mathbb{R}^{N_{\text{out}}}$, is defined as the function $\Phi_{\mathbf{W}}: \mathbb{R}^{N_{\text{in}}} \rightarrow \mathbb{R}^{N_{\text{out}}}$. For a given input $x \in \mathbb{R}^{N_{\text{in}}}$, the result $\Phi_{\mathbf{W}}(x) \in \mathbb{R}^{N_{\text{out}}}$ is computed as

$$\Phi_{\mathbf{W}}(x) := r_L(x),$$

where $r_L: \mathbb{R}^{N_{\text{in}}} \rightarrow \mathbb{R}^{N_{\text{out}}}$ is defined in a recursive manner using the functions $r_i: \mathbb{R}^{N_{\text{in}}} \rightarrow \mathbb{R}^{N_i}$ for $i = 0, \dots, L-1$, which are given by

$$\begin{aligned} r_L(x) &:= W_L r_{L-1}(x) + b_L, \\ r_i(x) &:= \rho_{N_i}^*(W_i r_{i-1}(x) + b_i) && \text{for } i = 1, \dots, L-1, \\ r_0(x) &:= x. \end{aligned}$$

Fitting neural network weights and biases to a given function f is accomplished by creating a sample set $T_{\text{train}} = \{(x_1, f(x_1)), \dots, (x_n, f(x_n))\} \subset X \times \mathbb{R}^{N_{\text{out}}}$ (the so-called *training set*), consisting of inputs $x_i \in X$ from an input set $X \subset \mathbb{R}^{N_{\text{in}}}$ and corresponding outputs $f(x_i) \in \mathbb{R}^{N_{\text{out}}}$. The process of finding the weights \mathbf{W} such that $\Phi_{\mathbf{W}}(x_i) \approx f(x_i)$ for $i = 1, \dots, n$ is called *training* of the neural network. During the training, the weights and biases of the neural network $\Phi_{\mathbf{W}}$ are iteratively adjusted such that a *loss function*, which measures the deviation of the output $\Phi_{\mathbf{W}}(x_i)$ for a given input x_i from the desired result $f(x_i)$, is minimized. A common choice for the loss function is the *mean squared error loss* $\mathcal{L}(\Phi_{\mathbf{W}}, T_{\text{train}})$ given as

$$\mathcal{L}(\Phi_{\mathbf{W}}, T_{\text{train}}) := \frac{1}{|T_{\text{train}}|} \sum_{(x,y) \in T_{\text{train}}} \|\Phi_{\mathbf{W}}(x) - y\|_2^2.$$

Popular examples of optimization algorithms used in neural network training are variants of (*stochastic gradient descent methods*, see [49] for an overview). For small neural networks with only a few layers and neurons, it is also possible to apply methods that use or approximate higher-order derivatives of the loss function, for instance, the *L-BFGS* optimizer [50], which is a limited-memory variant of the *BFGS* method, see for instance Section 6.1 in [43]. In the context of neural network training, each iteration of the optimizer is called *epoch*. Typically, a maximal number of epochs is prescribed for the optimizer to perform.

To prevent a neural network from *overfitting* the training data, we employ *early stopping* [51]. In this method, the loss function is evaluated on a *validation set* $T_{\text{val}} \subset X \times \mathbb{R}^{N_{\text{out}}}$ after each epoch. The validation set is usually chosen to be disjoint from the training set, i.e., $T_{\text{val}} \cap T_{\text{train}} = \emptyset$. Let \mathbf{W}_k denote the weights in epoch $k \in \mathbb{N}$. In each epoch, the value $\mathcal{L}(\Phi_{\mathbf{W}_k}, T_{\text{val}})$ is computed, and if this value does not decrease anymore over a prescribed number of consecutive epochs, the training is aborted. This method ensures that the resulting neural network can perform well on unseen data (that is assumed to have the same structure as the training data).

The result of the optimization routine typically depends strongly on the initial values \mathbf{W}_0 of the weights. There are several methods for initializing the weights of neural networks, for instance, the so-called *Kaiming initialization*, see [52] for more details. We perform *multiple restarts* of the training algorithm using different initial values for the weights to minimize the dependence of the resulting neural network on the weight initialization. Finally, we select the neural network $\Phi_{\mathbf{W}^*}$ that produced the smallest loss $\mathcal{L}(\Phi_{\mathbf{W}^*}, T_{\text{train}}) + \mathcal{L}(\Phi_{\mathbf{W}^*}, T_{\text{val}})$ over all training restarts, i.e., the smallest combined loss on the training and the validation set.

Finding an appropriate neural network architecture can be difficult in practical applications. Especially the number of layers and the number of neurons significantly influence the approximation capabilities of the resulting neural network. In addition, there are lots of different activation functions available. Typical examples include the *rectified linear unit (ReLU)* $\rho(x) = \max(x, 0)$, which is nowadays the most popular activation function [53], or the *hyperbolic tangent* $\rho(x) = \tanh(x) = \frac{e^{2x}-1}{e^{2x}+1}$.

5 Adaptive-ML-EnOpt algorithm using deep neural networks

The primary purpose of this work is to propose an adaptive machine learning-based algorithm for avoiding expensive FOM evaluations as often as possible. To this end, we first discuss the usage of DNNs for the NPV value and subsequently introduce the Adaptive-ML-EnOpt algorithm.

5.1 Surrogate models for the net present value

As discussed in Section 3, we use DNNs to construct a surrogate model for the FOM objective function J . DNNs are particularly well suited for non-intrusive model reduction if the simulator is considered a black box with no direct access to solutions of the underlying PDEs. In fact, given the formulation of the objective function (3b), we assume to only have access to the respective components $J_i(\mathbf{u})$.

Following the definition of a DNN in Section 4, two input-output maps can be used to approximate J . We refer to the *scalar-valued output* by considering $J: \mathbb{R}^{N_u} \rightarrow \mathbb{R}$ as the input-output map. Furthermore, we refer to the *vector-valued output* if we make different use of the structure of J by writing $J(\mathbf{u}) = \delta^T j(\mathbf{u})$ with

$$\begin{aligned} j &: \mathbb{R}^{N_u} \rightarrow \mathbb{R}^{N_t}, \\ j(\mathbf{u}) &:= [J_i(\mathbf{u})]_{i=1}^{N_t}, \end{aligned}$$

and the vector $\delta \in \mathbb{R}^{N_t}$, which includes the discount factors, is defined as

$$\delta := \left[\frac{1}{(1+d_t)^{\frac{t}{T}}} \right]_{i=1}^{N_t}.$$

In the scalar-valued case (DNN_s-approach), we directly construct a DNN for J with a corresponding function $\Phi_{\mathbf{w}_s}: \mathbb{R}^{N_u} \rightarrow \mathbb{R}$, i.e., we use a DNN with $N_{\text{in}} = N_u$ and $N_{\text{out}} = 1$. Instead, in the vector-valued case (DNN_v-approach), we construct a DNN for approximating j with a corresponding function $\Phi_{\mathbf{w}_v}: \mathbb{R}^{N_u} \rightarrow \mathbb{R}^{N_t}$ and, by using δ , we indirectly approximate J . This means that we apply a DNN with input and output dimensions given by $N_{\text{in}} = N_u$ and $N_{\text{out}} = N_t$, and multiply the result by δ whenever the respective DNN is used for approximating J . The algorithm described below works for both cases, the scalar-valued and the vector-valued output. Therefore, if access to the individual components of the vector-valued function j is available, it is possible to run the algorithm with both versions. The different neural network output sizes, and therefore, the various structures of the training data, might improve the DNN training results. In our numerical experiment, we observe that the vector-valued DNN yields slightly better results than the scalar-valued DNN (see Section 6). Nevertheless, we consider both the scalar- and vector-valued approaches to discuss the case where the black box reservoir simulator produces only scalar-valued outputs.

By the DNN_s- and DNN_v-approach, we thus construct a surrogate for the objective function for the optimization problem (3). It remains to explain a suitable and

robust EnOpt algorithm that takes advantage of a DNN but shows a similar convergence behavior as the FOM algorithm. A common strategy is to construct a sufficiently accurate surrogate $J_{\text{ML}} \in \{\Phi_{\mathbf{W}_s}, \delta^T \Phi_{\mathbf{W}_v}\}$ for the entire input space in a large offline time. Following the FOM-EnOpt procedure from Section 3, given J_{ML} , a surrogate-based procedure would then mean to set $F := J_{\text{ML}}$ in Algorithm 1. However, no FOM stopping criterion would be used, and since no error control for the surrogate model is given, no certification of the surrogate-based procedure would be available.

Remark 1 (Size of neural network input) Importantly, we observe that the input dimension N_u of both DNN approaches is proportional to the number of time steps N_t and the number of physical variables N_w in the model. Thus, dependent on the complexity of the reservoir simulation, N_u may be large. Consequently, it may not be possible to construct a surrogate model with a DNN that is accurate for the entire input space. Even if it were possible to construct such a DNN, we would require prohibitively costly training for computing the training set, validation set, and weights.

5.2 Adaptive algorithm

To circumvent the issue of constructing a globally accurate surrogate, in what follows, we describe the adaptive machine learning EnOpt algorithm (Adaptive-ML-EnOpt). In this algorithm, we incorporate the construction of the DNN into an outer optimization loop trained and certified by FOM quantities. With respect to the FOM-EnOpt procedure, we remark that each FOM optimization step requires N evaluations of J for computing \mathbf{d}_k . To obtain an appropriately accurate direction, it is required that N is chosen sufficiently large [54]. For the Adaptive-ML-EnOpt procedure, we only use a single FOM-based optimization step at each outer iteration k . Then, we use the N evaluations of the FOM as data points for training a locally accurate surrogate J_{ML}^k . Instead of proceeding with the FOM function J , we utilize the DNN to start an inner EnOpt algorithm with $F = J_{\text{ML}}^k$ as objective function in Section 3.1 and \mathbf{u}_k as initial guess. Denote by $\mathbf{u}_k^{(l)}$ the iterates of the inner optimization loop in the k -th outer iteration, i.e., in particular, we have $\mathbf{u}_k^{(0)} = \mathbf{u}_k$. According to (12), the inner EnOpt iteration terminates if the surrogate-based criterion

$$J_{\text{ML}}^k(\mathbf{u}_k^{(l)}) \leq J_{\text{ML}}^k(\mathbf{u}_k^{(l-1)}) + \varepsilon_i \quad (13)$$

is met for a suitable tolerance $\varepsilon_i > 0$. If the inner iteration terminates after L iterations with a control $\mathbf{u}_k^{(L)}$, the next outer iterate \mathbf{u}_{k+1} is defined as $\mathbf{u}_{k+1} := \mathbf{u}_k^{(L)}$. For a certified FOM-based stopping criterion of the outer optimization loop, given the iterate \mathbf{u}_k , we check whether the FOM-EnOpt procedure would, indeed, also stop at the same control point. Thus, we perform a single FOM-based optimization step, which includes the computation of \mathbf{d}_k and the line search, and results in a control $\tilde{\mathbf{u}}_k$. For verifying whether the FOM optimization step successfully finds a sufficiently increasing point at outer iteration k , we consider the FOM termination criterion

$$J(\tilde{\mathbf{u}}_k) \leq J(\mathbf{u}_k) + \varepsilon_o, \quad (14)$$

where $\varepsilon_o > 0$ is a suitable tolerance. If (14) is fulfilled, no improvement of the objective function value using FOM optimization steps can be expected, and therefore we also terminate the Adaptive-ML-EnOpt algorithm. If instead, (14) is not met, we use the computed training data (collected while computing \mathbf{d}_k) to retrain the DNN and restart an inner DNN-based EnOpt algorithm. We emphasize that the fully FOM-based stopping criterion constitutes a significant difference to what is proposed in [32], where the termination criterion is based on the approximation quality of the surrogate model at the current iterate. However, we saw in our experiments that such an approximation-based criterion might lead to an undesired early stopping of the algorithm.

One may be concerned about the fact that the surrogate-based inner optimization routine produces a decreasing or stationary point. For this reason, after every outer iteration k of the Adaptive-ML-EnOpt procedure, the inner DNN-optimization is only accepted after a sufficient increase, i.e.,

$$J(\mathbf{u}_{k+1}) > J(\mathbf{u}_k) + \varepsilon_o. \quad (15)$$

If an iterate is not accepted, we abort the algorithm. Instead of aborting, one may proceed with an intermediate FOM optimization step. We would further like to emphasize that the fulfillment of (15) also depends on the successful construction of the neural network, meaning that the parameters for the neural network are chosen appropriately. If, instead, (15) fails due to an inaccurate neural network, an automatic variation of the parameters could be enforced to the neural network training, and the corresponding outer iteration should be repeated. However, for the sake of simplicity and because it did not show any relevance in our numerical experiments, we do not specify approaches for the case that \mathbf{u}_{k+1} is not accepted due to (15).

While (15) prevents the algorithm from accepting an insufficient increase, the respective outer iteration point might still be far away from the optimum due to the lack of extrapolation accuracy. Since inner iterations are fast and the number of outer iterations is expected to be of moderate size, a longer optimization path compared to the reference (based solely on FOM evaluations) can still be considered computationally faster.

Remark 2 (Choice of the tolerances) Regarding the choice of the different tolerances ε_i and ε_o for the inner and outer stopping criteria in the Adaptive-ML-EnOpt algorithm, we propose to choose a small value for ε_i similar to the tolerance ε in the FOM-EnOpt procedure. The inner iterations are much cheaper due to the application of a fast surrogate, such that a more significant amount of inner iterations is acceptable. In contrast, we recommend selecting a larger tolerance ε_o to perform fewer outer iterations for obtaining a considerable speedup. However, if maximum convergence w.r.t. the FOM-EnOpt algorithm is desired, ε_o is to be set equal to ε .

The above-explained Adaptive-ML-EnOpt procedure is summarized in Algorithm 5 and additionally visualized in Fig. 1.

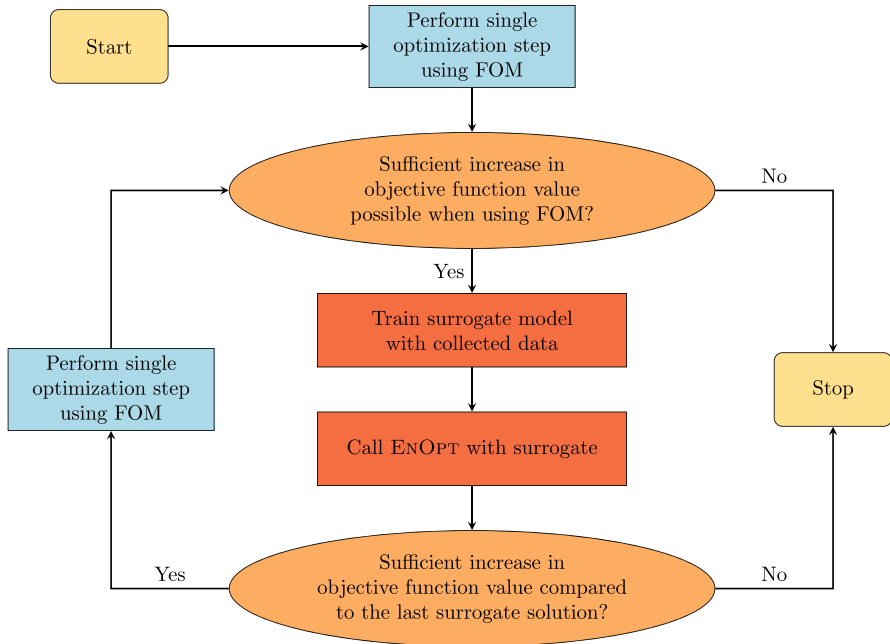


Fig. 1 Flowchart for the Adaptive-ML-EnOpt algorithm, cf. Algorithm 5

Algorithm 5 Adaptive-ML-EnOpt algorithm

Require: initial guess $\mathbf{u}_0 \in \mathbb{R}^{N_u}$, sample size $N \in \mathbb{N}$, tolerances $\varepsilon_o, \varepsilon_i > 0$ for outer and inner iterations, maximum number of outer and inner iterations k_o^*, k_i^* , DNN construction strategy $CS \in \{DNN_s, DNN_v\}$, DNN-specific variables V_{DNN} as discussed in Section 4 (e.g., network architecture, loss function, training parameters), initial step size $\beta > 0$, step size contraction $r \in (0, 1)$, maximum number of step size trials $v^* \in \mathbb{N}$

Ensure: approximate solution $\mathbf{u}^* \in \mathbb{N}^{N_u}$ of (3)

```

1: function ROMENOPT( $\mathbf{u}_0, N, \varepsilon_o, \varepsilon_i, k_o^*, k_i^*, CS, V_{DNN}, \beta, r, v^*$ )
2:    $\tilde{\mathbf{u}}_0, T_0 \leftarrow \text{OPTSTEP}[J]\mathbf{u}_0, N, 0, \beta, r, v^*$ 
3:    $k \leftarrow 0$ 
4:   while  $J(\tilde{\mathbf{u}}_k) > J(\mathbf{u}_k) + \varepsilon_o$  and  $k < k_o^*$  do
5:      $J_{ML}^k \leftarrow \text{TRAIN}T_k, CS, V_{DNN}$ 
6:      $\mathbf{u}_{k+1} \leftarrow \text{ENOPT}[J_{ML}^k]\mathbf{u}_k, N, \varepsilon_i, k_i^* \beta, r, v^*$ 
7:     if  $J(\mathbf{u}_{k+1}) \leq J(\mathbf{u}_k) + \varepsilon_o$  then
8:       return  $\mathbf{u}^* \leftarrow \mathbf{u}_k$ 
9:     end if
10:     $\tilde{\mathbf{u}}_{k+1}, T_{k+1} \leftarrow \text{OPTSTEP}[J]\mathbf{u}_{k+1}, N, k, \beta, r, v^*$ 
11:     $k \leftarrow k + 1$ 
12:  end while
13:  return  $\mathbf{u}^* \leftarrow \mathbf{u}_k$ 
14: end function
  
```

The TRAIN function performs the neural network training procedure as described in Section 4 and returns, depending on the chosen DNN construction strategy, a function Φ_{W_x} or Φ_{W_v} that approximates the FOM objective function J . Particularly, the result of TRAIN can be used as the function F in the ENOPT procedure.

The outer acceptance criterion (15) is checked in Line 7. Using the FOM-based stopping criterion in Line 4, we ensure that the Adaptive-ML-EnOpt algorithm has an equivalent stopping procedure as the FOM-EnOpt algorithm, see Line 4 in Algorithm 1. However, the algorithm might terminate at a different (local) optimal point, which we also observe in the numerical experiments.

Compared to the FOM-EnOpt procedure, we emphasize that, in the Adaptive-ML-EnOpt algorithm, mainly the single calls of OPTSTEP[J] in Lines 2 and 10 have FOM complexity, scaling with the number of samples N . Furthermore, the outer stopping criterion in Line 4 and the conditions for acceptance in Line 7 require a single FOM evaluation. The construction of the surrogate makes use of FOM data that is already available from the FOM optimization steps in Lines 2 and 10. In addition, while the training data for Line 5 is available from calling OPTSTEP[J], the training function TRAIN itself is relatively cheap. Furthermore, calling ENOPT[J_{ML}^k] has low computational effort since evaluating the surrogate J_{ML}^k for a given control (i.e., performing a single forward pass through the neural network) is much faster than evaluating J . The primary motivation for the Adaptive-ML-EnOpt algorithm is the idea that many of the costly FOM optimization steps in the FOM-EnOpt algorithm can be replaced by sequences of cheap calls of ENOPT[J_{ML}^k] with the surrogate J_{ML}^k . However, since the surrogate might only be reliable in a specific part of the set of feasible control vectors around the current iterate \mathbf{u}_k , we retrain the surrogate if the FOM optimization step suggests that a further improvement of the objective function value is possible. Therefore, the overall goal of the Adaptive-ML-EnOpt algorithm is to terminate with a considerably smaller number of (outer) iterations k than the FOM-EnOpt algorithm, and thus, to reduce the computational costs for

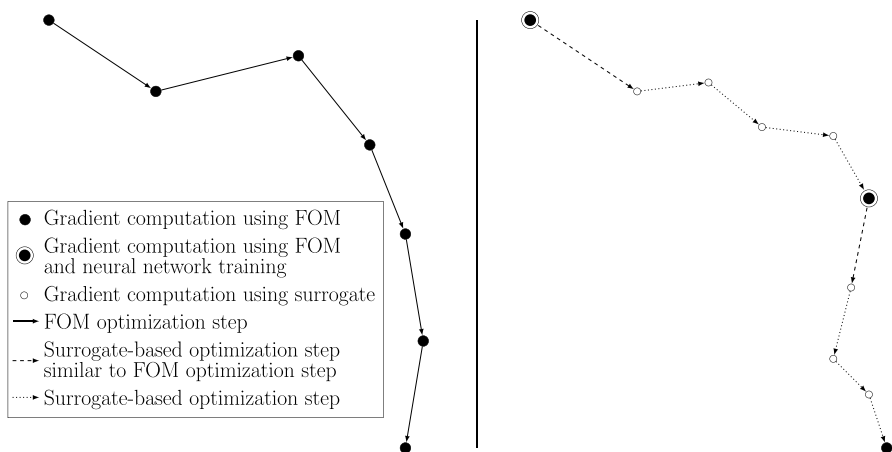


Fig. 2 Example of optimization paths taken by the FOM-EnOpt algorithm (left part of the figure) and the Adaptive-ML-EnOpt algorithm (right part of the figure)

solving the polymer EOR optimization problem in (3). We refer to the subsequent section for an extensive complexity and run time comparison in a practical example.

The main motivation for the Adaptive-ML-EnOpt algorithm is illustrated in Fig. 2. Computing the gradient information using evaluations of the function J is costly, whereas gradient computations using the approximation J_{ML}^k , obtained, for instance, via training a neural network, is cheap. In the example, the Adaptive-ML-EnOpt algorithm performs more optimization steps in total. However, most of these optimization steps are cheap since they only require evaluations of J_{ML}^k . For the Adaptive-ML-EnOpt algorithm, only those steps involving evaluations of J (i.e., outer iterations) require a large computational effort. Each optimization step is costly in the FOM-EnOpt algorithm since the exact objective function J is evaluated multiple times. Altogether, in the example shown in Fig. 2, the Adaptive-ML-EnOpt algorithm performs less costly gradient computations than the FOM-EnOpt procedure while arriving approximately at the same optimum. This motivates why the Adaptive-ML-EnOpt algorithm can be preferable with respect to the required computation time.

Remark 3 (Generalization of the method) The overall algorithmic procedure described in this section is applicable in many different scenarios and can be seen in a more general manner. For instance, the algorithm is independent of the exact objective function. Whenever the objective function is costly to evaluate and no certified surrogate is available, the Adaptive-ML-EnOpt might be useful. Hence, the algorithm can also be applied in scenarios where the objective function is not related to PDEs. Further, the choice of the surrogate is not limited to neural networks and machine learning algorithms in general. Any other approach that can be constructed from training data, similarly to DNNs, can serve as a replacement for the neural networks. In addition, the inner optimization loop is not restricted to the EnOpt algorithm. Whether or not gradient-based algorithms are applicable, depends on the surrogate and whether it is possible to compute its gradients with respect to the inputs.

6 Numerical validation for a five-spot benchmark problem

In this section, we present an example with a synthetic oil reservoir in which the polymer flooding optimization problem (3) is solved using the traditional solution method, the FOM-EnOpt algorithm, and our proposed Adaptive-ML-EnOpt method presented in Algorithm 5. The focus is to demonstrate a more efficient and improved method of dealing with the optimization part of a closed-loop reservoir workflow [41] for polymer flooding with the assumption that the geological properties of the reservoir are known. We start by providing information on the algorithm implementation.

6.1 Implementational details

For a numerical approximation of the system (1) of non-linear partial differential equations and the corresponding well equations (2), we make use of the open

porous media flow reservoir simulator (OPM) [35, 55]. The system is discretized spatially using a two-point flux approximation (TPFA) with upstream-mobility weighting (UMW) and temporally using a fully implicit Runge-Kutta method. The resulting discrete-in-time equations are solved using a Newton-Raphson scheme to obtain time-dependent states and the output quantities from the well's equation in terms of fluid production of the reservoir per time step. In this numerical experiment, we perform all polymer flooding simulations in parallel on a 50 core CPU.

For the implementation of the DNN-based surrogates, the Python package pyMOR [56] is used. The implementation of the neural networks and corresponding training algorithms in pyMOR is based on the machine learning library PyTorch [57].

Throughout our numerical experiments described in the subsequent section, we apply the L-BFGS optimizer with strong Wolfe line-search [58] for training the neural networks. Further, we perform a maximum of 1000 training epochs in each restart.

The number of training restarts influences the accuracy of the trained neural networks and the computation time required for the training. A larger number of restarts typically leads to smaller losses and more training time. To take these two factors into account, we consider different numbers of restarts in our numerical study presented below. The respective results can be found in the subsequent section. In general, we use relatively small numbers of restarts. First of all, we are not interested in obtaining a neural network with very high accuracy. Due to the adaptive retraining of the networks, the surrogates are replaced in each outer iteration anyway. They are only supposed to lead the optimizer to a point with a larger objective function value. On the other hand, as indicated before, a larger number of restarts might result in an unnecessarily long training phase, which must be performed in each outer iteration. The small numbers of 15 and 35 restarts we tried in our studies can thus be seen as a compromise between the accuracy of the surrogate models and computational effort for the training algorithm.

We use 10% of the sample set for validation during the neural network training, and the training routine is stopped early if the loss does not decrease for 10 consecutive epochs. Moreover, the mean squared error loss (MSE loss) is used as the loss function. The neural network training is performed on scaled data. The input values are scaled to $[0, 1]^{N_u}$, and the output values are scaled to $[0, 1]$ in the DNN_s -case and $[0, 1]^{N_v}$ in the DNN_v -case, respectively. The scaling of the input values can be computed exactly using the lower and upper bounds u_j^{low} and u_j^{upp} for the control variables by

$$u_j^i \mapsto \frac{u_j^i - u_j^{\text{low}}}{u_j^{\text{upp}} - u_j^{\text{low}}} \quad (16)$$

for $j = 1, \dots, N_w$ and $i = 1, \dots, N_t$. For the output values, we take the minimum and maximum value over the training set as lower and upper bounds and perform the same scaling as in Equation (16). The tanh function serves as the activation function for each layer. Kaiming initialization is applied for initializing the neural network weights.

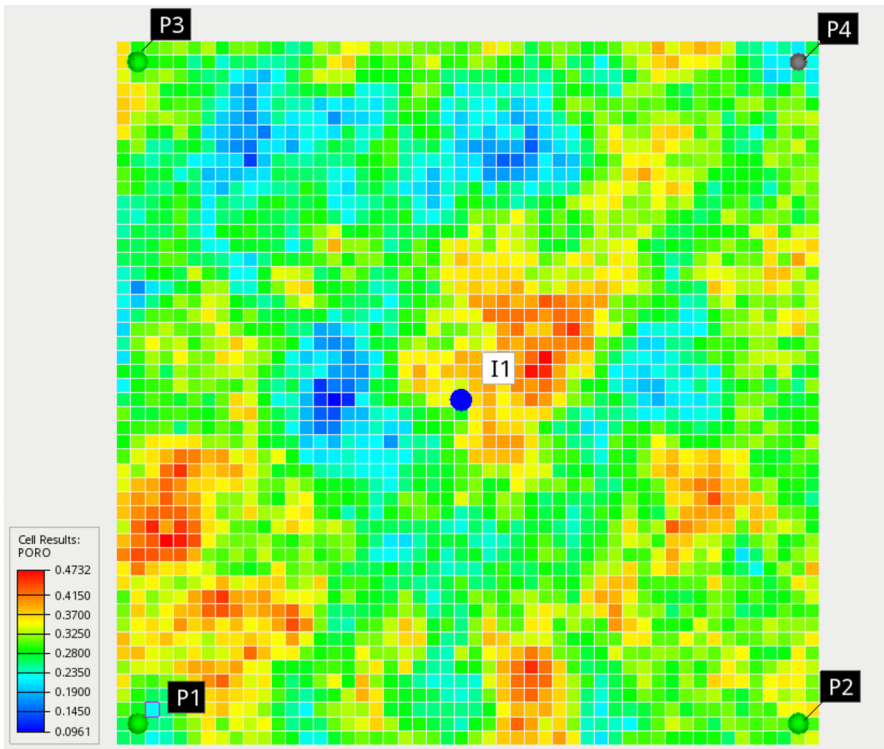


Fig. 3 Porosity distribution of the five-spot field and placement of the injection and production wells

The input and output dimensions of the neural networks were already described in Section 5 and are different for the DNN_s - and DNN_v -case. Regarding the training data for the vector-valued case DNN_v , we note that we require T_k to store $\{(\mathbf{u}_{k,j}, J(\mathbf{u}_{k,j}))\}_{j=1}^N$ instead of $\{(\mathbf{u}_{k,j}, J(\mathbf{u}_{k,j}))\}_{j=1}^N$, which we did not include in Algorithm 2 for brevity.

6.2 Case study: five-spot field

The numerical experiment considers a two-dimensional reservoir model with a three-phase flow, including oil, water, and gas (cf. Section 2). The computations are performed on a uniform grid that consists of 50×50 grid cells. The model has one injection and four production wells spatially arranged in a five-spot pattern as shown in Fig. 3.

On average, the reservoir has approximately 30% porosity with a heterogeneous permeability distribution. The initial reservoir pressure is 200 bar. The initial average oil and water saturations are 0.6546 and 0.3454, respectively. The original oil in place is $4.983 \cdot 10^6 \text{ sm}^3$. Fluid properties are similar to those of a light oil reservoir. The viscosity for saturated oil at varying bubble point pressure lies

Table 1 Reservoir model parameters used in the polymer flooding simulations

Parameter	Value	Unit
Dead pore space for polymer solution	0.1800	–
Maximum polymer adsorption value	$7.5 \cdot 10^{-4}$	kg/kg
Residual resistance factor of polymer solution	2.5	–
Reservoir rock density	1980	kg/rm ³
Polymer mixing parameters	0.65	–

between 0.1 cP and 0.56 cP, and the viscosity of water is 0.01 cP. The densities of oil and water are taken as 732 kg/m³ and 1000 kg/m³, respectively. In this setting, it is easy to see that the displacement is unfavorable since the oil-water mobility ratio λ is such that $10 \leq \lambda \leq 56$. The reservoir rock parameters utilized for the polymer flooding simulation in this problem are given by Table 1.

In this example, the injection well is controlled by two independent control variables, namely the water injection rate and the polymer concentration at each control time step. The lower and upper bounds for the water injection rate are set to 0 sm³/day and 2000 sm³/day respectively, while the lower and upper bounds for the polymer concentration are set to 0 kg/sm³ and 2.5 kg/sm³. Hence, the polymer injection rate ranges from 0 to 5000 kg/day. Each production well is controlled by a reservoir fluid production rate target with a lower limit of 0 sm³/day and an upper limit of 500 sm³/day. Bottom hole pressure limits are imposed on the wells, namely a maximum of 500 bar for the injector and a minimum of 150 bar for each producer. The production period for the reservoir is set to 50 months, and the control time step is taken as 5 months. Therefore, there are $N_u = (2 + 4) \times 10 = 60$ control variables in total to solve for in (3). For the objective function (3a), we used the economic parameters listed in Table 2.

Using the two different surrogate models for the objective function (3b) constructed by means of neural networks, namely DNN_s and DNN_v , as explained in Section 5, the optimization problem (3) is solved using the Adaptive-ML-EnOpt algorithm. In this case, the Adaptive-ML-EnOpt algorithm for (3) using DNN_s and DNN_v to approximate the objective function J from (3b) is denoted by AML-EnOpt_s and AML-EnOpt_v, respectively. The EnOpt parameters for both, the FOM-EnOpt and the two variants of the Adaptive-ML-EnOpt method, are

Table 2 Economic parameters used in the numerical experiments

Parameter	Value	Unit
Oil price r_{OP}	500	USD/sm ³
Price of gas production r_{GP}	0.15	USD/sm ³
Cost of polymer injection r_{PI}	2.5	USD/kg
Cost of polymer production r_{PP}	0.5	USD/kg
Cost of water injection or production r_{WP}, r_{WP}	30	USD/sm ³
Annual discount rate d_t	0.1	–

Table 3 Parameters used in the FOM-EnOpt and Adaptive-ML-EnOpt algorithms

Parameter	Value
Initial constant oil rate in \mathbf{u}_0^1 and \mathbf{u}_0^2	150 and 100 (sm ³ /day)
Initial constant water rate in \mathbf{u}_0^1 and \mathbf{u}_0^2	700 and 600 (sm ³)
Initial constant polymer concentration in \mathbf{u}_0^1 and \mathbf{u}_0^2	0.5 and 0.5 (kg/sm ³)
Initial step size β_0	0.3
Step size contraction r	0.5
Maximum step size trials ν^*	10
Initial control-type variance σ_j	0.001
Constant correlation factor ρ	0.9
Perturbation size N	100
FOM-EnOpt ε	10^{-6}
Tolerances Adaptive-ML-EnOpt inner iteration ε_i	10^{-6}
Adaptive-ML-EnOpt outer iteration ε_o	10^{-2}

presented in Table 3. We note that the tolerances ε , ε_i , and ε_o are chosen as suggested in Remark 2 and are applied to the scaled quantities, i.e., the output quantities, for which the respective stopping criteria in Algorithm 1 are checked, have already been scaled as described in Section 6.1.

We compare the Adaptive-ML-EnOpt results with those of the FOM-EnOpt algorithm for two different initial guesses $\mathbf{u}_0^1 \in \mathcal{D}_{\text{ad}}$ and $\mathbf{u}_0^2 \in \mathcal{D}_{\text{ad}}$. The initial solution \mathbf{u}_0^1 includes 700 sm³/day for the water injection rate at the injection well, 150 sm³/day for the reservoir fluid production rate at each production well, and 0.5 kg/sm³ for the polymer concentration (equivalently 350 kg/day for polymer injection rate) at the injection well over the simulation period. Similarly, \mathbf{u}_0^2 includes 600 sm³/day for the water injection rate, 100 sm³/day for the reservoir fluid production rate, and 0.5 kg/sm³ for the polymer concentration.

Figure 4 compares the values of the objective function during the outer iterations of the FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_v strategies using the initial solutions \mathbf{u}_0^1 and \mathbf{u}_0^2 . Furthermore, the value $J(\mathbf{u}_k)$ at the outer iterate \mathbf{u}_k (denoted by “FOM value”) for the respective Adaptive-ML-EnOpt method is depicted.

Since the Adaptive-ML-EnOpt algorithms only use an approximate surrogate model J_{ML}^k , the values of J and J_{ML}^k are not necessarily the same for the control \mathbf{u}_k . This behavior is especially apparent in Fig. 4b, where the AML-EnOpt_v algorithm is examined for the initial guess \mathbf{u}_0^1 . Here, after the first outer iteration, the values $J(\mathbf{u}_1)$ and $J_{\text{ML}}^0(\mathbf{u}_1)$ differ from each other by a significant amount. A possible reason is that the surrogate model J_{ML}^0 does not extrapolate well to the region where the first (inner) Adaptive-ML-EnOpt iteration converged to. This further indicates that the found iterate \mathbf{u}_1 is far from the initial solution \mathbf{u}_0 , where the initial model J_{ML}^0 was trained. However, since the Adaptive-ML-EnOpt algorithm uses evaluations of J in the stopping criterion, the Adaptive-ML-EnOpt does not terminate but continues by training a new surrogate model using training data sampled normally around \mathbf{u}_1 . Hence, the new surrogate J_{ML}^1 tries to approximate the objective function J well

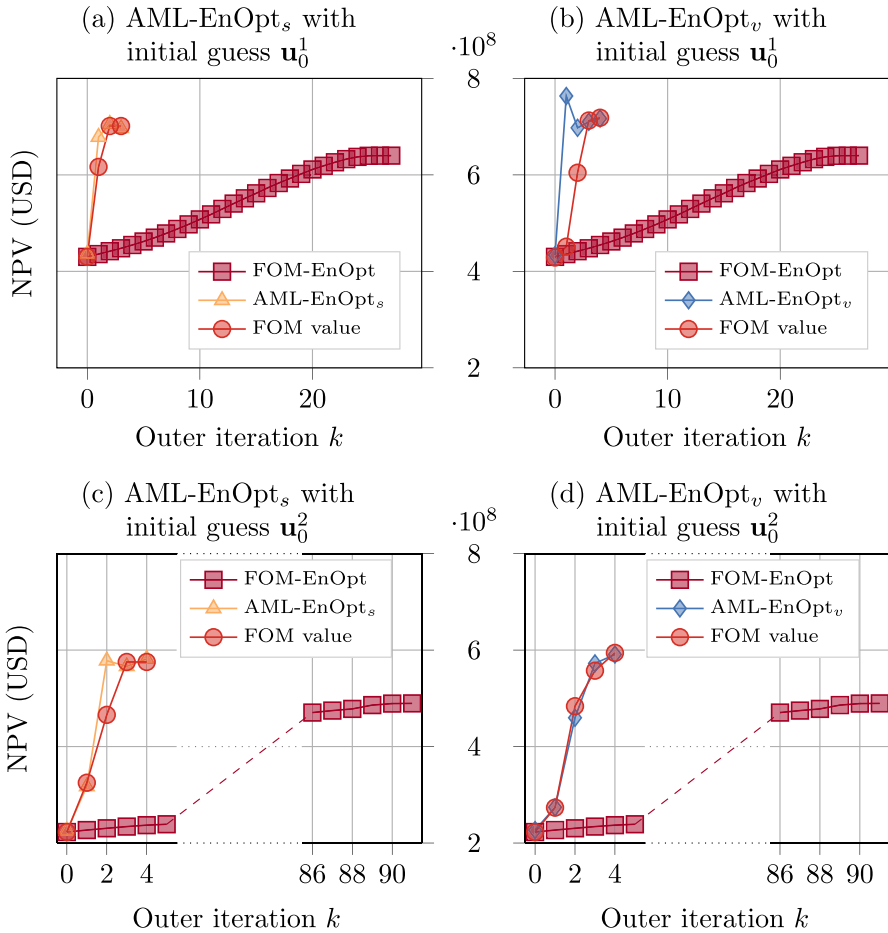


Fig. 4 Comparison of the NPV values obtained during the outer iterations of the FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_v procedures for two different initial guesses $\mathbf{u}_0^1 \in \mathcal{D}_{ad}$ and $\mathbf{u}_0^2 \in \mathcal{D}_{ad}$. For each Adaptive-ML-EnOpt procedure, the corresponding FOM value $J(\mathbf{u}_k)$ at the current iterate \mathbf{u}_k of the respective Adaptive-ML-EnOpt method is indicated as well

around \mathbf{u}_1 . In each plot, we see that in the last two outer iterations of the respective Adaptive-ML-EnOpt procedure, the FOM value and the Adaptive-ML-EnOpt value agree to minimal deviations. This suggests that the surrogate model approximates the full objective function well in the region of the (local) optimum found by the Adaptive-ML-EnOpt method.

More so, in Fig. 4, it is seen that both, the AML-EnOpt_s and the AML-EnOpt_v, algorithm, require considerably less (costly) outer iterations than the FOM-EnOpt method. This leads to an improvement in the run time of the method, which is detailed in Table 4. Besides the faster convergence of the method, we also remark that the Adaptive-ML-EnOpt algorithms find local optima with larger objective

Table 4 Comparisons of the results from the different solution strategies FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_t, using the initial guess \mathbf{u}_0^1 and $N_1 = N_2 = 35$ neurons in each hidden layer and 15 restarts for the neural network training

Method	FOM value	Surrogate value	Outer iter.	Inner iter.	FOM eval.	Surrogate eval.	Training time (min)	T_{total} (min)	Speedup
FOM-EnOpt	$6.400 \cdot 10^8$	—	28	—	2839	—	—	54.86	—
AML-EnOpt _s	$7.013 \cdot 10^8$	$6.968 \cdot 10^8$	4	233	407	12315	2.03	8.87	6.18
AML-EnOpt _t	$7.185 \cdot 10^8$	$7.168 \cdot 10^8$	5	312	509	14101	3.40	14.10	3.89

Table 5 Comparisons of the results from the different solution strategies FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_t, using the initial guess u_0^2 and $N_1 = N_2 = 25$ neurons in each hidden layer and 35 restarts for the neural network training

Method	FOM value	Surrogate value	Outer iter.	Inner iter.	FOM eval.	Surrogate eval.	Training time (min)	T_{total} (min)	Speedup
FOM-EnOpt	$4.895 \cdot 10^8$	–	92	–	9310	–	–	135.77	–
AML-EnOpt _s	$5.754 \cdot 10^8$	$5.816 \cdot 10^8$	4	111	407	10837	2.89	9.85	13.78
AML-EnOpt _t	$5.942 \cdot 10^8$	$5.908 \cdot 10^8$	4	117	407	10033	4.40	11.05	12.29

function values than the FOM-EnOpt algorithm. However, since the objective function J is multi-modal, this is not guaranteed.

We emphasize that each outer iteration of the Adaptive-ML-EnOpt algorithm includes many inner iterations (see also Tables 4 and 5), which leads to the large jumps in the objective function values between consecutive outer iterations, as presented in Fig. 4.

Further comparisons in terms of function values, numbers of inner and outer iterations, numbers of evaluations of the FOM function J and surrogate approximations J_{ML}^k , training times of the neural networks, total run time (including the training times), and speedup are presented in Tables 4 and 5.

With the different initial guesses \mathbf{u}_0^1 and \mathbf{u}_0^2 , we found that the number of outer iterations required by the FOM-EnOpt algorithm significantly differs. However, the Adaptive-ML-EnOpt methods require only 4 and 5 outer iterations. This reduced number of outer iterations leads to a remarkable speedup in the overall computation time T_{total} and is particularly reflected in the reduced number of FOM evaluations, i.e., evaluations of the objective function J , which require costly polymer flooding simulations. Although each outer iteration consists of multiple inner iterations using the surrogate J_{ML}^k , it does not contribute substantially to the overall run time because evaluating the surrogate J_{ML}^k is very cheap.

For the initial solution \mathbf{u}_0^1 , the optimizers obtained from the three solution strategies are depicted in Fig. 5. Further, the initial guess \mathbf{u}_0^1 is shown as a reference.

The control variables obtained by the AML-EnOpt_s and the AML-EnOpt_v algorithm are close to those of the FOM-EnOpt method, except for production well 3 (see Fig. 5c) and the water injection rate (see Fig. 5e). For each control variable, the values obtained via the AML-EnOpt_s and AML-EnOpt_v procedures are close to each other. Together with the FOM values of AML-EnOpt_s and AML-EnOpt_v, presented in Table 4 and the evolution of the FOM values for the two methods shown in Fig. 4a–b, this suggests that the AML-EnOpt_s and the AML-EnOpt_v methods traverse almost the same path in the control space \mathcal{D}_{ad} and find local optima close to each other.

Figure 6a depicts a comparison of the total field oil production for the optimal solutions (in Fig. 5) of the three solution methods. The total field oil production by FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_v are $1.343 \cdot 10^6$, $1.425 \cdot 10^6$, and $1.554 \cdot 10^6$ (in sm^3), respectively. The solution obtained by AML-EnOpt_v attains the highest oil production in total, followed by the AML-EnOpt_s. The total back-produced water and polymer from operating the five-spot field with the different optimal solutions are presented in Fig. 6b and c, respectively. Here, we found that the AML-EnOpt_s and AML-EnOpt_v solutions are more economical and environmentally friendly than the one provided by using the FOM-EnOpt method.

To further investigate the effects of different neural network architectures on the resulting NPV values, Fig. 7 depicts the NPV values obtained by the Adaptive-ML-EnOpt algorithm when using different numbers of neurons in the hidden layers of the surrogate models DNN_s and DNN_v .

We observe that the AML-EnOpt_v method results are very similar, which suggests that the DNN_v -approach is more robust and leads to similar optimal solutions independent of the neural network structure. In the case of the AML-EnOpt_s algorithm, different numbers

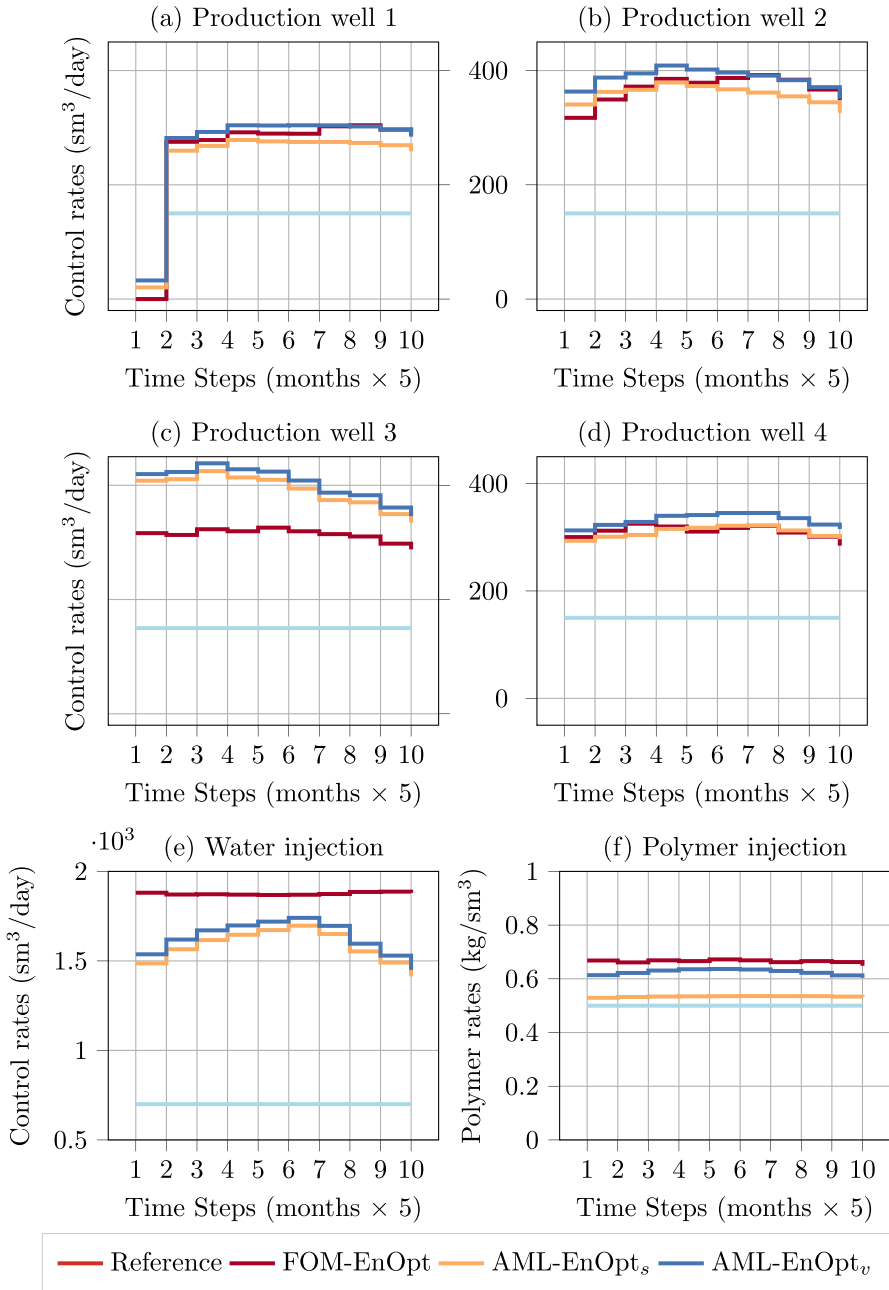


Fig. 5 Comparison of the optimal solutions obtained via the FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_v algorithms using the initial guess u_0^1 , which is depicted as the reference solution

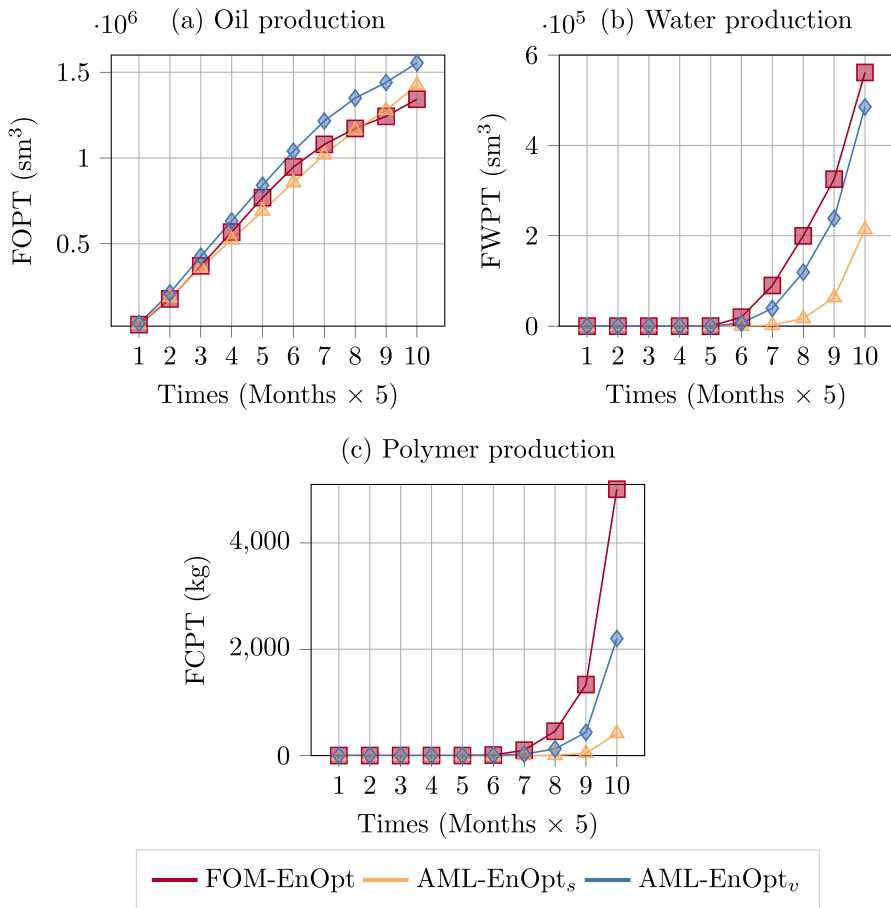


Fig. 6 Comparison of the production data obtained from the different solution strategies FOM-EnOpt, AML-EnOpt_s, and AML-EnOpt_v, using the initial guess \mathbf{u}_0^1

of neurons lead to results with a larger variation. In particular, the number of outer iterations performed is different. Hence, the architecture of the underlying network seems to have a significant effect on the performance of the resulting AML-EnOpt_s algorithm.

The maximum, minimum, and average training and validation losses that occurred in the AML-EnOpt_s and AML-EnOpt_v algorithm for the initial guess \mathbf{u}_0^1 are presented in Table 6. The table shows the respective MSE losses for different numbers of neurons in the hidden layers.

The results in Table 6 do not suggest a significant influence of the number of neurons on the training and validation results. Further, the scalar- and vector-valued cases, DNN_s and DNN_v , respectively, perform similarly in overall training and validation losses. However, we emphasize that, in the DNN_v case, the MSE loss cannot be related directly to the difference in the output function. Instead, one has to take into account that the outputs of DNN_v are summed up to obtain the surrogate J_{ML}^k , while the MSE loss is measured on the vector-valued outputs of the neural network.

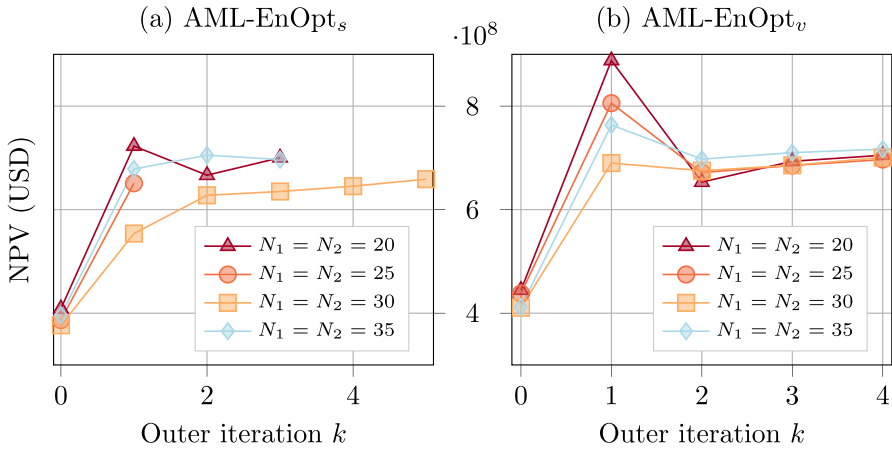


Fig. 7 Comparison of the Adaptive-ML-EnOpt procedures AML-EnOpt_s and AML-EnOpt_v for different numbers of neurons in the hidden layers with fixed initial guess \mathbf{u}_0^1

Table 6 Maximum, minimum, and average MSE loss in the AML-EnOpt_s and AML-EnOpt_v algorithm with different numbers of neurons in the hidden layers of the neural networks DNN_s and DNN_v for fixed initial guess \mathbf{u}_0^1 . The number of hidden layers is fixed to two

Method	Neurons	Outer iter.	Training loss			Validation loss		
			Max.	Min.	Avg.	Max.	Min.	Avg.
DNN _s	20	4	$1.2 \cdot 10^{-4}$	$1.3 \cdot 10^{-6}$	$5.3 \cdot 10^{-5}$	$5.4 \cdot 10^{-3}$	$6.7 \cdot 10^{-5}$	$2.3 \cdot 10^{-3}$
DNN _s	25	2	$6.0 \cdot 10^{-4}$	$1.3 \cdot 10^{-6}$	$3.0 \cdot 10^{-4}$	$2.3 \cdot 10^{-3}$	$2.1 \cdot 10^{-3}$	$2.2 \cdot 10^{-3}$
DNN _s	30	6	$7.9 \cdot 10^{-4}$	$8.5 \cdot 10^{-7}$	$1.7 \cdot 10^{-4}$	$6.6 \cdot 10^{-3}$	$1.7 \cdot 10^{-3}$	$3.6 \cdot 10^{-3}$
DNN _s	35	4	$1.8 \cdot 10^{-4}$	$6.1 \cdot 10^{-6}$	$8.2 \cdot 10^{-5}$	$5.5 \cdot 10^{-3}$	$3.7 \cdot 10^{-4}$	$2.7 \cdot 10^{-3}$
DNN _v	20	5	$1.8 \cdot 10^{-3}$	$2.1 \cdot 10^{-5}$	$5.2 \cdot 10^{-4}$	$6.9 \cdot 10^{-3}$	$1.2 \cdot 10^{-3}$	$4.2 \cdot 10^{-3}$
DNN _v	25	5	$9.9 \cdot 10^{-4}$	$1.5 \cdot 10^{-5}$	$4.1 \cdot 10^{-4}$	$6.4 \cdot 10^{-3}$	$9.4 \cdot 10^{-4}$	$3.6 \cdot 10^{-3}$
DNN _v	30	5	$9.0 \cdot 10^{-4}$	$9.9 \cdot 10^{-6}$	$4.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-2}$	$5.2 \cdot 10^{-4}$	$4.3 \cdot 10^{-3}$
DNN _v	35	5	$6.0 \cdot 10^{-4}$	$1.1 \cdot 10^{-6}$	$2.2 \cdot 10^{-4}$	$8.8 \cdot 10^{-3}$	$4.6 \cdot 10^{-4}$	$4.3 \cdot 10^{-3}$

Altogether, the numerical experiments with different numbers of neurons suggest that already small DNNs with only 20 neurons in each of the hidden layers yield appropriate results. In this specific application, we do not benefit from increasing the complexity of the neural network. We have seen the same behavior when using more than two hidden layers.

7 Conclusion and future work

In this contribution, we presented a new algorithm to speed up PDE-constrained optimization problems occurring in the context of enhanced oil recovery. The algorithm is based on adaptively constructed surrogate models that make use of deep neural networks for approximating the objective function. In each outer iteration of

the algorithm, a new surrogate model is trained with data consisting of full-order function evaluations around the current control point. Afterwards, an ensemble-based optimization algorithm is applied to the surrogate to obtain a candidate for the next iteration. We perform full order model evaluations to validate whether the resulting controls correspond to a local optimum of the true objective function. These function evaluations also serve as training data for constructing the next surrogate.

Our numerical results confirm that the described algorithm can accelerate the solution of the enhanced oil recovery optimization problem. At the same time, in our numerical experiments, the procedure produces controls with even larger objective function values than those obtained using only costly full-order model evaluations. However, we should emphasize that such an improvement in the objective function value is not guaranteed and, in our case, results from the multi-modal structure of the objective function.

The investigated five-spot benchmark problem served as a proof of concept for our Adaptive-ML-EnOpt algorithm, where FOM evaluations were relatively quickly accessible, and the input dimension was of moderate size. Future research is thus devoted to more involved numerical experiments with more significant complexity.

As indicated in the optimization problem description, we focused on a scenario with fixed geological properties. However, in practical applications, these geological parameters are usually unknown and typically treated by ensemble-based methods, where the ensemble is to be understood not only with respect to perturbations of the controls for approximating the gradient but also with respect to different samples of geological properties. One of the central future research perspectives is incorporating such geological uncertainty in our algorithm. The main challenge is the high dimension of the space of possible geological parameters. Naively using these parameters as additional inputs for the neural network is thus not feasible. Future research might consider reducing the dimension of the space of geological parameters by incorporating additional information on the distribution of such parameters and passing the reduced variables to the neural networks.

Furthermore, replacing neural networks as surrogate models for the objective function, for instance, by polynomial approximations obtained via linear regression or by different machine learning approaches, such as kernel methods [59], could be investigated further. The Adaptive-ML-EnOpt algorithm is formulated in such a way that replacing the surrogate model and its training is readily possible. Any approximation of the objective function built from evaluations of the true objective function is feasible and can directly be used in the algorithm. In addition, the inner iterations are not restricted to the EnOpt procedure but can also be performed using different optimization routines. However, we should emphasize that in the current formulation, no information on the exact gradient, neither of the true objective function nor the surrogate model, is required. This might change when employing different optimization routines. Moreover, the presented approach is not restricted to the NPV objective function in enhanced oil recovery but can be generalized to any scalar-valued quantity of interest. The algorithm might be of particular relevance in cases where no direct access to the underlying PDE solutions is possible, and no error estimation for the surrogate model is available.

Finally, we note that the selection of parameters, such as the initial and adaptive sample covariance for the approximate gradient computation of the objective function, can influence the behavior of the EnOpt method. This impact has been examined in greater detail by [54]. It will be of future interest to see how the different selections of these parameters can impact the final solutions of AML-EnOpt₁ and AML-EnOpt₂.

Funding Open access funding provided by University Of Stavanger • Tim Keil, Hendrik Kleikamp, Micheal Oguntola and Mario Ohlberger received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2044 –390685587, Mathematics Münster: Dynamics–Geometry–Structure.

• Tim Keil and Mario Ohlberger received funding from the Deutsche Forschungsgemeinschaft under contract OH 98/11-1.

• Micheal Oguntola and Rolf Lorentzen received funding from the Research Council of Norway and the industry partners, ConocoPhillips Skandinavia AS, Aker BP ASA, Vår Energi AS, Equinor Energy AS, Neptune Energy Norge AS, Lundin Energy Norway AS, Halliburton AS, Schlumberger Norge AS, and Wintershall Dea Norge AS, of The National IOR Centre of Norway.

Declarations

Conflict of interest The authors declare no competing interests.

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. Pancholi, S., Negi, G.S., Agarwal, J.R., Bera, A., Shah, M.: Experimental and simulation studies for optimization of water–alternating–gas (CO₂) flooding for enhanced oil recovery. *Petroleum Res.* **5**(3), 227–234 (2020). <https://doi.org/10.1016/j.ptlrs.2020.04.004>
2. Van, S.L., Chon, B.H.: Well-pattern investigation and selection by surfactant-polymer flooding performance in heterogeneous reservoir consisting of interbedded low-permeability layer. *Kor. J. Chem. Eng.* **33** (12), 3456–3464 (2016). <https://doi.org/10.1007/s11814-016-0190-7>
3. Zhang, Y., Lu, R., Forouzanfar, F., Reynolds, A.C.: Well placement and control optimization for WAG/SAG processes using ensemble-based method. *Comput. Chem. Eng.* **101**, 193–209 (2017). <https://doi.org/10.1016/j.compchemeng.2017.02.020>
4. Gudiña, E.J., Fernandes, E.C., Rodrigues, A.I., Teixeira, J.A., Rodrigues, L.R.: Biosurfactant production by bacillus subtilis using corn steep liquor as culture medium. *Front. Microbiol.* **6**, 59 (2015). <https://doi.org/10.3389/fmicb.2015.00059>
5. Abidin, A., Puspasari, T., Nugroho, W.: Polymers for enhanced oil recovery technology. *Procedia Chemistry* **4**, 11–16 (2012). <https://doi.org/10.1016/j.proche.2012.06.002>
6. Wang, D., Seright, R.S., Shao, Z., Wang, J., et al.: Key aspects of project design for polymer flooding at the daqing oilfield. *SPE Reserv Eval Eng* **11**(06), 1–117 (2008). <https://doi.org/10.2118/109682-PA>


7. Oguntola, M.B., Lorentzen, R.J.: On the robust value quantification of polymer eor injection strategies for better decision making. In: ECMOR XVII. <https://doi.org/10.3997/2214-4609.202035057>. European Association of Geoscientists & Engineers, vol. **2020**, pp 1–25 (2020)
8. Xu, L., Zhao, H., Li, Y., Cao, L., Xie, X., Zhang, X., Li, Y.: Production optimization of polymer flooding using improved monte carlo gradient approximation algorithm with constraints. *J. Circ. Syst. Comput.* **27** (11), 1850167 (2018). <https://doi.org/10.1142/S0218126618501670>
9. Zhou, K., Hou, J., Zhang, X., Du, Q., Kang, X., Jiang, S.: Optimal control of polymer flooding based on simultaneous perturbation stochastic approximation method guided by finite difference gradient. *Comput. Chem. Eng.* **55**, 40–49 (2013). <https://doi.org/10.1016/j.compchemeng.2013.04.009>
10. Dedè, L.: Reduced basis method and error estimation for parametrized optimal control problems with control constraints. *J. Sci. Comput.* **50**(2), 287–305 (2012)
11. Grepl, M.A., Kärcher, M.: Reduced basis a posteriori error bounds for parametrized linear-quadratic elliptic optimal control problems. *C.R. Math.* **349**(15-16), 873–877 (2011)
12. Kammann, E., Tröltzsch, F., Volkwein, S.: A posteriori error estimation for semilinear parabolic optimal control problems with application to model reduction by POD. *ESAIM: M2AN* **47**(2), 555–581 (2013). <https://doi.org/10.1051/m2an/2012037>
13. Kärcher, M., Tokoutsis, Z., Grepl, M.A., Veroy, K.: Certified reduced basis methods for parametrized elliptic optimal control problems with distributed controls. *J. Sci. Comput.* **75**(1), 276–307 (2018)
14. Negri, F., Rozza, G., Manzoni, A., Quateroni, A.: Reduced basis method for parametrized elliptic optimal control problems. *SIAM J. Sci. Comput.* **35**(5), 2316–2340 (2013). <https://doi.org/10.1137/120894737>
15. Benner, P., Ohlberger, M., Patera, A., Rozza, G., Urban, K. (eds.): Model reduction of parametrized systems. MS&A. modeling, simulation and applications, vol. 17. Springer, Berlin (2017). <https://doi.org/10.1007/978-3-319-58786-8>. Selected papers from the 3rd MoRePaS Conference held at the International School for Advanced Studies (SISSA), Trieste, October 13–16, 2015
16. Garmatter, D., Haasdonk, B., Harrach, B.: A reduced basis Landweber method for nonlinear inverse problems. *Inverse Prob.* **32**(3), 035001–21 (2016). <https://doi.org/10.1088/0266-5611/32/3/035001>
17. Zahr, M.J., Farhat, C.: Progressive construction of a parametric reduced-order model for PDE-constrained optimization. *Int. J. Numer. Meth. Engng* **102**, 1111–1135 (2015). <https://doi.org/10.1002/nme.4770>
18. Keil, T., Mechelli, L., Ohlberger, M., Schindler, F., Volkwein, S.: A non-conforming dual approach for adaptive trust-region reduced basis approximation of PDE-constrained parameter optimization. *ESAIM Math. Model. Numer. Anal.* **55**(3), 1239–1269 (2021). <https://doi.org/10.1051/m2an/2021019>
19. Qian, E., Grepl, M., Veroy, K., Willcox, K.: A certified trust region reduced basis approach to PDE-constrained optimization. *SIAM J. Sci. Comput.* **39**(5), 434–460 (2017). <https://doi.org/10.1137/16M1081981>
20. Yue, Y., Meerbergen, K.: Accelerating optimization of parametric linear systems by model order reduction. *SIAM J. Optim.* **23**(2), 1344–1370 (2013). <https://doi.org/10.1137/120869171>
21. Gavrilenko, P., Haasdonk, B., Iliev, O., Ohlberger, M., Schindler, F., Toktaliev, P., Wenzel, T., Youssef, M.: A full order, reduced order and machine learning model pipeline for efficient prediction of reactive flows. arXiv:2104.02800v2 (2021)
22. Haasdonk, B., Kleikamp, H., Ohlberger, M., Schindler, F., Wenzel, T.: A new certified hierarchical and adaptive RB-ML-ROM surrogate model for parametrized PDEs. arXiv:2204.13454 (2022)
23. Haasdonk, B., Ohlberger, M., Schindler, F.: An adaptive model hierarchy for data-augmented training of kernel models for reactive flow. arXiv:2110.12388v1 (2021)
24. Hesthaven, J.S., Ubbiali, S.: Non-intrusive reduced order modeling of nonlinear problems using neural networks. *J. Comput. Phys.* **363**, 55–78 (2018). <https://doi.org/10.1016/j.jcp.2018.02.037>
25. Wang, Q., Hesthaven, J.S., Ray, D.: Non-intrusive reduced order modeling of unsteady flows using artificial neural networks with application to a combustion problem. *J. Comput. Phys.* **384**, 289–307 (2019). <https://doi.org/10.1016/j.jcp.2019.01.031>
26. Lee, K., Carlberg, K.T.: Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *J. Comput. Phys.* **404**, 108973–32 (2020). <https://doi.org/10.1016/j.jcp.2019.108973>
27. Ahmadi, M.A.: Developing a robust surrogate model of chemical flooding based on the artificial neural network for enhanced oil recovery implications. *Math. Probl. Eng.* **2015**, 9 (2015). <https://doi.org/10.1155/2015/706897>

28. Saberi, H., Esmailnezhad, E., Choi, H.J.: Artificial neural network to forecast enhanced oil recovery using hydrolyzed polyacrylamide in sandstone and carbonate reservoirs. *Polymers* **13**(16). <https://doi.org/10.3390/polym13162606> (2021)
29. Cheraghi, Y., Kord, S., Mashayekhizadeh, V.: Application of machine learning techniques for selecting the most suitable enhanced oil recovery method; challenges and opportunities. *J. Pet. Sci. Eng.* **205**, 108761 (2021). <https://doi.org/10.1016/j.petrol.2021.108761>
30. Golzari, A., Haghghat Sefat, M., Jamshidi, S.: Development of an adaptive surrogate model for production optimization. *J. Pet. Sci. Eng.* **133**, 677–688 (2015). <https://doi.org/10.1016/j.petrol.2015.07.012>
31. Lee, J.-Y., Shin, H.-J., Lim, J.-S.: Selection and evaluation of enhanced oil recovery method using artificial neural network. *Geosystem Engineering* **14**, 157–164 (2011). <https://doi.org/10.1080/12269328.2011.10541345>
32. Lye, K.O., Mishra, S., Ray, D., Chandrashekar, P.: Iterative surrogate model optimization (ISMO): An active learning algorithm for PDE constrained optimization with deep neural networks. *Comput. Methods Appl. Mech. Eng.* **374**, 113575 (2021). <https://doi.org/10.1016/j.cma.2020.113575>
33. Wang, S., Bhouiri, M.A., Perdikaris, P.: Fast PDE-constrained optimization via self-supervised operator learning. *arXiv:2110.13297* (2021)
34. Sarma, P., Durlafsky, L.J., Aziz, K., Chen, W.H.: Efficient real-time reservoir management using adjoint-based optimal control and model updating. *Comput. Geosci.* **10** (1), 3–36 (2006). <https://doi.org/10.1007/s10596-005-9009-z>
35. Rasmussen, A.F., Sandve, T.H., Bao, K., Lauser, A., Hove, J., Skaflestad, B., Klöfkorn, R., Blatt, M., Rustad, A.B., Sævreid, O., et al: The open porous media flow reservoir simulator. *Computers & Mathematics with Applications* **81**, 159–185 (2021). <https://doi.org/10.1016/j.camwa.2020.05.014>
36. Bao, K., Lie, K. -A., Møyner, O., Liu, M.: Fully implicit simulation of polymer flooding with MRST. *Comput. Geosci.* **21**(5), 1219–1244 (2017). <https://doi.org/10.1007/s10596-017-9624-5>
37. Holmes, J.: Enhancements to the strongly coupled, fully implicit well model: wellbore crossflow modeling and collective well control. In: *SPE Reservoir Simulation Symposium*. <https://doi.org/10.2118/12259-MS.OnePetro> (1983)
38. Holmes, J., Barkve, T., Lund, O.: Application of a multisegment well model to simulate flow in advanced wells. In: *European Petroleum Conference*. <https://doi.org/10.2118/50646-MS.OnePetro> (1998)
39. Chen, Z.: *Reservoir simulation: Mathematical techniques in oil recovery*. SIAM (2007)
40. Lu, R., Reynolds, A.: Joint optimization of well locations, types, drilling order, and controls given a set of potential drilling paths. *SPE J.* **25** (03), 1285–1306 (2020). <https://doi.org/10.2118/193885-MS>
41. Chen, Y., Oliver, D.S., Zhang, D.: Efficient ensemble-based closed-loop production optimization. *SPE J.* **14**(04), 634–645 (2009). <https://doi.org/10.2118/112873-PA>
42. 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>
43. Nocedal, J., Wright, S.: *Numerical optimization* springer. <https://doi.org/10.1007/978-0-387-40065-5> (2006)
44. Montgomery, D.C., Jennings, C.L., Kulahci, M.: *Introduction to time series analysis and forecasting*. Wiley, Hoboken (2015)
45. Fonseca, R.R.-M., Chen, B., Jansen, J.D., Reynolds, A.: A stochastic simplex approximate gradient (stosag) for optimization under uncertainty. *Int. J. Numer. Methods Eng.* **109**(13), 1756–1776 (2017). <https://doi.org/10.1002/nme.5342>
46. Oguntola, M.B., Lorentzen, R.J.: Ensemble-based constrained optimization using an exterior penalty method. *J. Pet. Sci. Eng.* **207**, 109165 (2021). <https://doi.org/10.1016/j.petrol.2021.109165>
47. Petersen, P., Voigtlaender, F.: Optimal approximation of piecewise smooth functions using deep reLU neural networks. *Neural Netw.* **108**, 296–330 (2018). <https://doi.org/10.1016/j.neunet.2018.08.019>
48. Elbrächter, D., Grohs, P., Jentzen, A., Schwab, C.: DNN Expression rate analysis of high-dimensional PDEs: Application to option pricing. *Constructive Approximation*. <https://doi.org/10.1007/s00365-021-09541-6> (2021)
49. Bottou, L., Curtis, F.E., Nocedal, J.: Optimization methods for large-scale machine learning. *SIAM Rev.* **60**(2), 223–311 (2018). <https://doi.org/10.1137/16M1080173>

50. Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. *Math. Program.* **45**, 503–528 (1989). <https://doi.org/10.1007/BF01589116>
51. Prechelt, L.: Early stopping - but when?. In: *Neural Networks: Tricks of the Trade*, Volume 1524 of LNCS, Chapter 2. https://doi.org/10.1007/978-3-642-35289-8_5, pp 55–69. Springer (1997)
52. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 1026–1034. <https://doi.org/10.1109/ICCV.2015.123> (2015)
53. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–444 (2015). <https://doi.org/10.1038/nature14539>
54. Fonseca, R., Kahrobaei, S., Van Gastel, L., Leeuwenburgh, O., Jansen, J.: Quantification of the impact of ensemble size on the quality of an ensemble gradient using principles of hypothesis testing. In: *SPE Reservoir Simulation Symposium*. <https://doi.org/SPE-173236-MS>. OnePetro (2015)
55. Baxendale, D., Rasmussen, A.F., Rustad, A.B., Skille, T., Sandve, T.H.: *OPM Flow documentation manual*. Open Porous Media Initiative (2021)
56. Milk, R., Rave, S., Schindler, F.: PyMOR – generic algorithms and interfaces for model order reduction. *SIAM J. Sci. Comput.* **38**(5), 194–216 (2016). <https://doi.org/10.1137/15m1026614>
57. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: *Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R (eds.) Advances in Neural Information Processing Systems*, vol. 32, pp 8024–8035. Curran Associates Inc (2019)
58. Wolfe, P.: Convergence conditions for ascent methods. *SIAM Rev.* **13**(2), 185–188 (1971). <https://doi.org/10.1137/1013035>
59. Hofmann, T., Schölkopf, B., Smola, A.J.: Kernel methods in machine learning. *Ann. Stat.* **36**(3), 1171–1220 (2008). <https://doi.org/10.1214/009053607000000677>

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

Authors and Affiliations

Tim Keil¹ · Hendrik Kleikamp¹ · Rolf J. Lorentzen² · Micheal B. Oguntola^{2,3}  · Mario Ohlberger¹

Tim Keil
tim.keil@uni-muenster.de

Hendrik Kleikamp
hendrik.kleikamp@uni-muenster.de

Rolf J. Lorentzen
rolo@norceersearch.no

Mario Ohlberger
mario.ohlberger@uni-muenster.de

¹ Institute for Analysis and Numerics and Mathematics Münster, University of Münster, Einsteinstrasse 62, D-48149 Münster, Germany

² NORCE-Norwegian Research Center AS, 5838 Bergen, Norway

³ University of Stavanger, 4036 Stavanger, Norway