

Closed-loop workflow for short-term optimization of wind-powered reservoir management

Mathias M. Nilsen^{a,b}^{*}, Rolf J. Lorentzen^a, Olwijn Leeuwenburgh^c, Andreas S. Stordal^a, Eduardo Barros^c

^a NORCE RESEARCH AS, Nygårdsgaten 112, 5008 Bergen, Norway

^b Department of Mathematics, University of Bergen, Allégaten 41, 5007 Bergen, Norway

^c TNO, Princetonlaan 6, Utrecht, Netherlands

ARTICLE INFO

Keywords:

Reservoir management
Multi-scale optimization
Wind power
Closed-loop

ABSTRACT

This paper presents a closed-loop workflow for short-term optimization of reservoir management powered by offshore wind energy. Motivated by the need to reduce CO₂ emissions in the Norwegian oil and gas sector, the workflow integrates wind power forecasts into optimization of daily well control while adhering to a long-term production strategy optimized for economic output. The workflow utilizes coarsened reservoir models calibrated through ensemble-based data assimilation to minimize computational costs. A realistic benchmark model, Drogon, is used to demonstrate the methodology. The workflow dynamically adjusts well rates to align power demand with wind availability, minimizing reliance on gas turbines and thereby reducing emissions. The numerical experiment demonstrates the potential of the workflow by significantly reducing short-term emissions without compromising the NPV.

1. Introduction

The Norwegian offshore oil and gas industry is bound by regulatory obligations to reduce its emissions significantly over the next decade. As a result, considerable attention has been given to the potential use of power from nearby offshore wind farms to operate oil and gas fields. Traditionally, offshore oil and gas production relies on gas turbines, which are by far the largest contributors to emissions on the Norwegian continental shelf (Offshore Norge, 2024). Wind power, in contrast, generates energy with little to no emissions. However, the high variability and uncertainty of wind power present several challenges from a reservoir management perspective.

Offshore oil and gas production requires substantial power, often reaching several tens of megawatts. Some key processes that drive this high power demand include water injection for pressure support and for guiding oil and gas towards the production wells, as well as compression of produced gas for reinjection or export (gas sales). The power demand of both of these processes depends on the chosen well rates and pressures (well controls) of the reservoir operation. Typically, well rates are selected or optimized based on the economic performance of the reservoir's production. However, several studies have included CO₂ emissions in the well control optimization to study the trade-off between profits and emissions. Angga et al. (2022) studied how

an increasing CO₂ tax affects the rates in waterflooding optimization. They concluded that including a CO₂ term in the objective resulted in a more energy-efficient waterflooding strategy. Both Oliver et al. (2025) and Oliver et al. (2024) use the weighted sum approach for multi-objective optimization to obtain a Pareto front, where the two objectives are profit and emissions. In Oliver et al. (2024), the optimization was performed on a real field case from the North Sea. They were able to reduce the projected emissions significantly compared to the operator base case with minimal reductions in profit. In Nilsen et al. (2025), a multi-objective optimization was performed over an ensemble of realistic long-term time series for daily available wind power on a benchmark reservoir. They concluded that including wind power makes it possible to reduce emissions significantly (as high as 40%) without too much loss in revenue.

In this paper, we develop a novel closed-loop workflow that, in each iteration of the loop, incorporates a wind power forecast to reduce CO₂ emissions by aligning the short-term power demand of operation with the forecasted wind power, while still tracking a long-term target operational strategy. The long-term operational strategy is based on maximizing the life-cycle net present value (NPV), whereas the short-term emission reduction strategy is based on minimizing emissions while penalizing deviations from the long-term target. The key novelty

^{*} Corresponding author at: NORCE RESEARCH AS, Nygårdsgaten 112, 5008 Bergen, Norway.

E-mail address: mani@norce-research.no (M.M. Nilsen).

Table 1
Schedule of when wells are added to Drogon.

	A1	A2	A5	A3	A4	A6	OP5
Type	Producer	Producer	Injector	Producer	Producer	Injector	Producer
Date	05.01.2018	10.03.2018	08.05.2018	13.07.2018	22.09.2018	17.11.2018	01.07.2020

of the proposed approach lies in the joint conditioning of daily well rate optimization on both wind power forecasts and a predefined long-term production strategy, which results in lower emissions from reservoir production without compromising long-term economic output.

To reduce the computational cost of the workflow, each short-term optimization is performed with a coarsened model of the full reservoir model. The coarsened model is recalibrated at the beginning of each time window before the optimization. This is inspired by the work of Sandve et al. (2024), which showed the potential of using coarsened reservoir models for optimization in reservoir management. In this paper, we use a realistic benchmark reservoir model called Drogon (Equinor, 2025a) to demonstrate the closed-loop workflow. Since this study focuses on the variability of wind power, it is assumed that for each time window, a reliable forecast of daily average wind speed (and, therefore, wind power) exists. For simplicity, the forecast uncertainty is neglected in this paper.

The proposed workflow of this paper is inspired by the work of van Essen et al. (2013), which introduced a two-level strategy that used a model predictive controller to track a long-term operational strategy for oil and gas reservoir management. Additionally, the workflow developed in this paper shares some similarity with the work of Huo et al. (2023), where they developed a bi-level optimization framework that aligns renewable energy generation with the energy-intensive demands of aluminum electrolysis. Their study demonstrates how a strategic configuration of wind, solar, and energy storage can reduce emissions without significantly compromising production capacity.

The outline of the paper is the following. In Section 2, the Drogon reservoir model is described. This section also discusses how the coarse reservoir model is created and recalibrated using data assimilation. Section 3 describes how the power demand and CO₂ emission from the production are modeled. Section 4 describes how the long-term strategy is determined, while Section 5 describes the proposed workflow of the paper. The numerical experiments and results are described in Section 6. The paper ends with a final discussion and outlook in Section 7.

2. The Drogon reservoir

In this paper, we employ the Drogon reservoir, developed by Equinor (2025a). The Drogon field has a size of approximately $3.6 \times 4.5 \times 0.1$ kilometers, and is divided into $43 \times 73 \times 31$ cells, of which 70972 are active. Although small, the field is designed to mimic realistic reservoir behavior and contains several faults, as shown in Fig. 1, which depicts the grid-geometry of the Drogon reservoir and its production and injection wells. An ensemble of a hundred different realizations of the geological reservoir properties is created. However, for simplicity, only one realization corresponding to the mean of the ensemble is used in this paper. This reduces the number of simulations we need to run to demonstrate the proposed workflow. The mean porosity (PORO) and permeability in the x-direction (PERMX) are depicted in Figs. 2(a) and 3(a).

The simulation of Drogon is divided into two periods: one for history matching and one for prediction and optimization. The historical period starts on January 1, 2018, and runs until July 1, 2020. During the first year of this period, the wells labeled A1 to A6 are added sequentially (schedule shown in Table 1). At the beginning of the prediction period, an additional production well (labeled OP5) is added. This period lasts 4.5 years until January 1, 2025. The workflow proposed in this paper is demonstrated on this prediction period.

2.1. The coarsened drogon model

One approach to make a computationally fast proxy model for short-term optimization is grid coarsening, which reduces the number of cells in the numerical reservoir simulator. This is an alternative approach to other popular methods like physics-informed neural networks (see Liu et al., 2025 for an overview of physics-informed neural networks). In this paper, the Python package *pycopm* (Landa-Marbán, 2024) is used for coarsening the Drogon reservoir model. The grid is coarsened by a factor of two in each dimension, x, y, and z, such that the grid size goes from $43 \times 73 \times 31$ to $23 \times 37 \times 16$. The resulting upscaled porosity and permeability are depicted in Figs. 2(b) and 3(b). The coarsened model is generally not sufficiently accurate to be directly used in the optimization workflow. However, the accuracy can be improved significantly by the use of history matching (also referred to as data assimilation).

2.2. History matching

The term history matching refers to calibrating a model to historical data by minimizing the difference between the data and the predictions of the model. However, in this paper, we instead calibrate the coarsened model to match the fine-scale model via simulation data generated by the fine-scale model. Sandve et al. (2024) did this for the Drogon model, where they looked at different levels of coarsening and recalibrated the relative permeability curves of different flow regions. A simpler approach is taken in this paper, where the cell values for porosity and permeability of the coarse model are recalibrated instead.

This paper uses ensemble-based data assimilation to recalibrate the coarse model to the fine-scale model. Ensemble-based data assimilation (Evensen et al., 2022) is derived from Bayes' theorem, which is given by

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}, \quad (1)$$

where $p(\theta|y)$ is called the posterior density of the parameters, θ , conditioned on the data y . The posterior density is proportional to the product of the likelihood, $p(y|\theta)$, and the prior density, $p(\theta)$. Most ensemble-based data assimilation methods seek to obtain an ensemble of realizations, θ_n^{post} (for $n = 1$ to N), from the posterior density by updating an ensemble, θ_n^{pr} , drawn from the prior density. A common approach is to assume a Gaussian density for the likelihood and prior density. The task of updating the prior ensemble is then reduced to minimizing an ensemble of cost functions on the form (see Evensen et al., 2022 for details):

$$\mathcal{L}_n(\theta) = \frac{1}{2}(\theta - \theta_n^{\text{pr}})^\top \Sigma_\theta^{-1}(\theta - \theta_n^{\text{pr}}) + \frac{1}{2}(\mathcal{M}_C(\theta) - y_n)^\top \Sigma_y^{-1}(\mathcal{M}_C(\theta) - y_n), \quad (2)$$

for $n = 1$ to N and simulated data given by $y_n = \mathcal{M}_F(\theta_F) + \epsilon_n$. Here, $\mathcal{M}_C(\theta)$ is the simulation output from the coarse model given the parameters, θ , and $\mathcal{M}_F(\theta_F)$ is the output of the fine-scale model (given the fine-scale parameters, θ_F). Further, $\epsilon_n \sim \mathcal{N}(0, \Sigma_y)$, and Σ_y is the error covariance matrix of the data, which in our case comes from the simulation of the fine-scale model. In the case where y contains real data, Σ_y represents the measurement error. However, when we calibrate a coarse model to a fine-scale model, we do not have measurement errors. As the coarse model cannot reproduce the simulated data from the fine-scale model, we still include ϵ_n as a type of

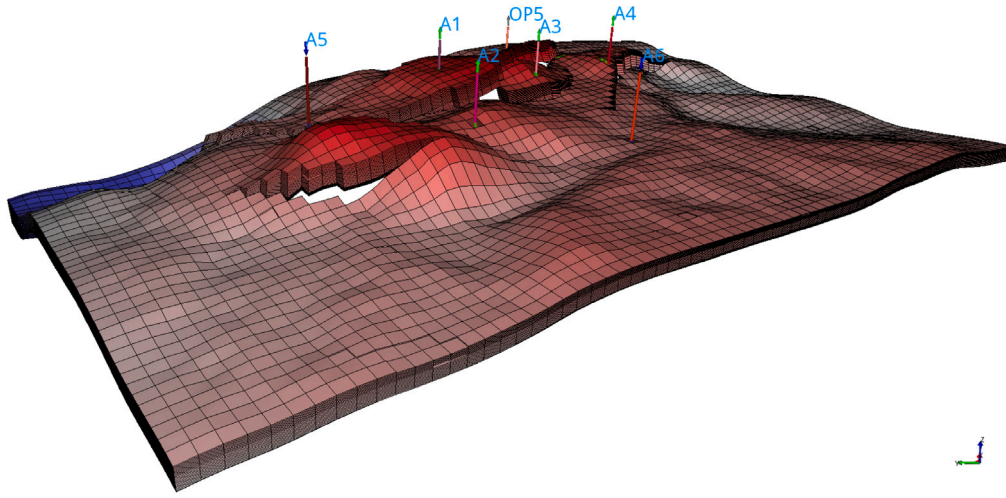


Fig. 1. Grid structure of the Drogon reservoir. The color represents depth.

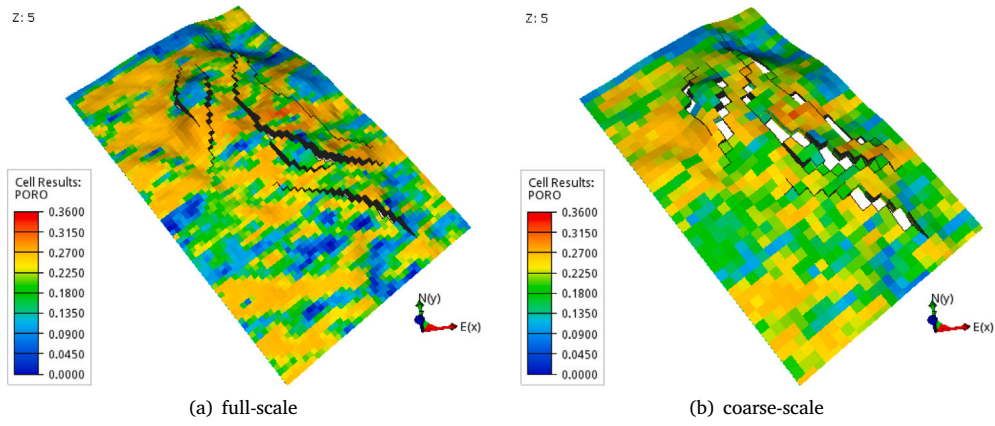


Fig. 2. Porosity (PORO).

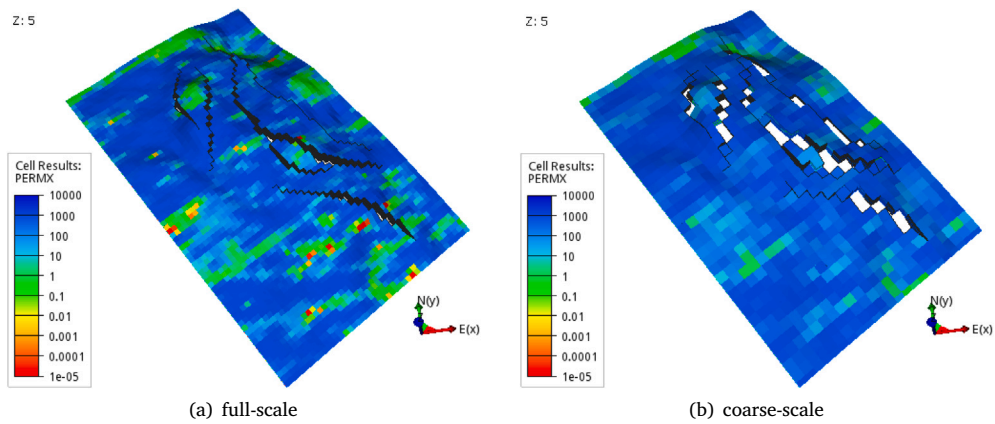


Fig. 3. Permeability in x-direction (PERMX).

regularization. Lastly, the covariance for the prior distribution is given by Σ_θ .

To illustrate how history matching works, we recalibrate the permeability and porosity of the coarse model for the history period of the Drogon field. The output data of the simulation are the oil (WOPR), gas (WGPR), and water (WWPR) production rates of the wells A1 to A4.

The simulation output of the fine-scale model is shown as the dashed line in Fig. 4. The injection rate in the injectors A5 and A6 is set to 8000 Sm³/day. The prior ensemble is generated by sampling a hundred different realizations of a Gaussian field with spatial correlation, where the mean is given by the coarsened porosity and permeability from

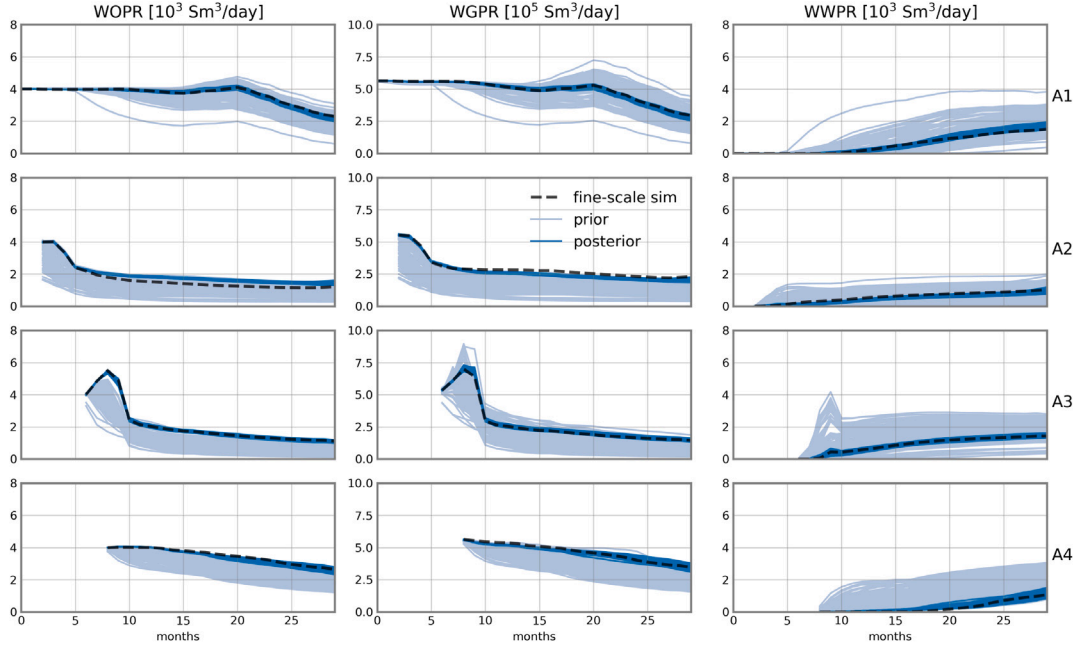


Fig. 4. The oil, gas, and water production rates for the prior (light-blue) and posterior (blue) ensemble in the production wells during the history period. The black dashed lines are the data from the fine-scale model simulation.

Figs. 2(b) and 3(b). The light-blue lines in Fig. 4 show the prior ensemble simulation output, while the dark-blue lines depict the posterior ensemble output.

3. Modeling power demand and CO₂ emissions

When modeling the production's topside facility, four different energy consumers are included: a gas compressor for gas export, a water injection pump, a water treatment system for produced water, and lastly, a constant baseload of 4 MW. To cover the total power demand, we assume the availability of two nearby 8 MW offshore wind turbines. In periods when wind power is insufficient, gas turbines installed on the platform will provide the necessary backup power.

3.1. Gas compression

Following the recipe from My Engineering Tools (2025), the power demand of gas compression can be calculated by

$$P_{\text{comp}} = \frac{nZRT_1}{(n-1)M\eta_{\text{poly}}} \left[\left(\frac{p_2}{p_1} \right)^{\frac{n-1}{n}} - 1 \right] Q_m^g, \quad (3)$$

where Q_m^g is the mass flow rate of gas given in kg/s. R is the gas constant, M is the molecular weight, Z is the gas compressibility factor (assumed to have the value one here), p_1 and p_2 are the inlet and outlet pressures, and T_1 is the temperature at the inlet. Gas typically needs to be compressed from around 3 to 20 bar to around 70 to 100 bar (Rasmussen and Kurz, 2009) before it can be sent to export pipes. Therefore, we use $p_1 = 5$ bar and $p_2 = 100$ bar. η_{poly} is the polytropic efficiency of the compression, and n is the polytropic coefficient, which is given by

$$n = \frac{1}{1 - (k-1)/(k \cdot \eta_{\text{poly}})}, \quad (4)$$

where $k = 1.32$ is the heat capacity ratio of natural gas. The polytropic efficiency is determined by the volumetric flow rate, Q_v^g , and can be approximated (Guo et al., 2007) by

$$\eta_{\text{poly}}(Q_v^g) = 0.61 + 0.03 \log(0.5885 Q_v^g), \quad (5)$$

when Q_v^g is given in Sm³/h. In this paper, we assume a gas density at standard conditions of 0.7215 kg/Sm³. This number is based on the energy content of a standard cubic meter of natural gas, which is 11.111 kWh (Norwegian Petroleum, 2025), and the specific energy content of natural gas, which is $S_e = 15.4$ kWh/kg (The Engineering ToolBox, 2025). The power demand and efficiency of the compression are shown in Fig. 5 as a function of gas rate.

3.2. Water pumping and treatment

The power demand of a water pumping system can be calculated by the following (Angga et al., 2022):

$$P_{\text{pump}} = \frac{\rho g H Q_v^w}{\eta}, \quad (6)$$

where $g = 9.81$ m/s² is the gravitational acceleration and $\rho = 1025$ kg/m³ is the density of salt water. Q_v^w is the volumetric flow rate of water. H is the head (in meters), which is given by

$$H = \frac{p_2 - p_1}{\rho g}, \quad (7)$$

where p_1 is the suction pressure and p_2 is the discharge pressure. For simplicity, we take the suction pressure to be 1 bar, and the discharge pressure to be the largest top head pressure of the different injection wells (Angga et al., 2022). We assume that our pumping system consists of only one pump with an efficiency curve (model taken from the appendix of Angga et al., 2022) shown in the right panel of Fig. 6. The left panel of Fig. 6 shows the power demand as a function of pumping rate when $p_2 = 200$ bar.

For produced water, we assume that the energy used for treating water is 1.358 J/Sm³ (Angga et al., 2022), which is approximately 3.772 kWh/Sm³.

3.3. Power generation and emissions

To cover the total power demand, denoted P_{tot} , given by

$$P_{\text{tot}} = P_{\text{comp}} + P_{\text{pump}} + P_{\text{water-treatment}} + P_{\text{baseload}}, \quad (8)$$

it is assumed that two 8 MW offshore wind turbines are available. The power curve used for the wind turbines is displayed in the left panel




Fig. 5. (left panel) Power demand as a function of the rate for the gas compression. (right panel) Polytopic efficiency as a function of the rate.




Fig. 6. (left panel) Power demand of injection pump as a function of the pump rate. (right panel) Pump efficiency as a function of the rate.




Fig. 7. (left panel) Model of the power curve of an 8 MW wind turbine. (right panel) Daily average wind speed (blue line) for four months, and daily available wind power (brown line).

of Fig. 7, and is found in Appendix D in Musial et al. (2016). The blue line in the right panel of Fig. 7 shows daily average wind speed for four months from a relevant location in the North Sea (latitude of 61.3°N and longitude of 2.3°E) taken from the NORA10EI dataset (Haakenstad et al., 2020). The NORA10EI dataset is an open-access hindcast for the North Sea, the Norwegian Sea and the Barents Sea for the period 1979 to 2017. The brown line illustrates the available wind power calculated from the power curve in the left panel.

It is assumed that whenever the wind turbines are not able to cover the total power demand, gas turbines will cover the remaining load. The power demand that the gas turbines need to cover on any given day is given by

$$P_{\text{gas}} = \max(P_{\text{tot}} - P_{\text{wind}}, 0), \quad (9)$$

where P_{wind} is the available wind power. The number of gas turbines needed to cover P_{gas} is determined by

$$n_{\text{gas}} = \left\lceil \frac{P_{\text{gas}}}{P_{\text{max}}} \right\rceil, \quad (10)$$

where $\lceil \cdot \rceil$ is the ceiling function and $P_{\text{max}} = 14.9 \text{ MW}$ is the maximum power output of a single gas turbine. We assume that the load P_{gas} is equally distributed among the gas turbines, such that each turbine's load is $P_{\text{turbine}} = P_{\text{gas}}/n_{\text{gas}}$. The efficiency curve of a single gas turbine is shown in the right panel of Fig. 8 (this model is taken from Angga

et al., 2022). The gas-fuel rate (mass rate) of the gas turbine system is then given by

$$Q_{\text{fuel}} = n_{\text{gas}} \frac{P_{\text{turbine}}}{\eta S_e}, \quad (11)$$

where $\eta = \eta(P_{\text{turbine}})$ is the efficiency (right panel of Fig. 8). The gas fuel demand is subtracted from the gas produced by the reservoir. The CO_2 emission rate is then given by

$$Q_{\text{CO}_2} = S_{\text{CO}_2} Q_{\text{fuel}}, \quad (12)$$

where $S_{\text{CO}_2} = 2.75 \text{ kg CO}_2/\text{kg fuel}$ (The Engineering ToolBox, 2025). The left panel of Fig. 8 shows the CO_2 emission rate as a function of the gas turbine system power load.

4. Long-term operational strategy

The workflow proposed in this paper starts with a target (long-term) operational strategy for the well controls. In principle, this strategy can be calculated based on any desired long-term outcome. In this paper, the target control strategy is determined by maximizing the long-term economic output of the production, i.e., maximizing the NPV. NPV is defined as the discounted sum of revenues minus costs, and is




Fig. 8. (left panel) CO₂ emission rate as a function of the turbine system power load. (right panel) Efficiency curve of a single gas turbine.

mathematically defined in this paper as

$$NPV = \sum_{i=1}^{N_i} \frac{c_o V_o^i + c_g V_g^i - c_{wi} V_{wi}^i - c_{wp} V_{wp}^i - c_{co2} M_{co2}^i}{(1+r)^{t_i/365}}, \quad (13)$$

where c_o and c_g are the prices of oil and gas in units of USD/Sm³. V_o^i and V_g^i are the volumes of oil and gas produced in time interval i . The quantities c_{wi} and c_{wp} are the cost of water injection and treatment, respectively. Further, V_{wi}^i and V_{wp}^i are the volumes of injected and produced water in time interval i . c_{co2} is the price of emission (CO₂ tax) and M_{co2}^i is the mass of emitted CO₂ in time interval i . r is the annual discount rate and t_i is the number days since $i = 0$.

The price of oil is set to 75 USD per barrel (approximately 472 USD per Sm³), while the price of gas is set to 0.4 USD per Sm³ (based on Equinor's average internal gas price in 2023 and 2024 [Equinor, 2025b](#)). The cost of water injection and treatment is set to 4 and 6 USD per barrel (approximately 25.2 and 37.7 USD per Sm³), respectively. The combined cost of CO₂ tax and quotas for the Norwegian oil and gas industry is projected to reach 2000 NOK per tonne CO₂ (approximately 188 USD per tonne) by the year 2030 ([Offshore Norge, 2024](#)), we therefore set $c_{co2} = 150$ USD per tonne.

The target controls (to be determined by optimization) are oil production rates in each producer and water injection rates in each injector of the Drogon reservoir. The target controls are allowed to change every six months during the 4.5 years of the optimization period, which gives a total number of 63 (7 wells × 9 changes) controls to determine. The well rates are subject to upper bounds and a minimum rate of 0. The maximum oil production rate in wells A2-A4 is 4000 Sm³/day, and the maximum for the production wells A1 and OP5 is 3000 Sm³/day. The maximum water injection rates in the two injectors, A5 and A6, are 8000 Sm³/day. The optimization of the target controls is determined based on the full-scale model and with the assumption that we always have enough wind power to cover the power demand. The initial controls of the optimization are set to the resulting well rates from a reference simulation where the target oil production rate in each producer is set to the maximum, and the injection rate in the injectors is set to 6500 Sm³. For the reference simulation, a production well is closed if the water cut in that well reaches a value higher than 88% as this is the threshold for uneconomical water production ([Fonseca et al., 2020](#)). The initial controls are shown as red lines in [Fig. 9](#). The gray lines are the optimized target controls, and the blue lines are the actual oil production rates resulting from the target controls (gray lines).

The total oil, gas, and water production rates from the target strategy and the reference simulation are shown in the upper and middle rows in [Fig. 10](#). The total power demand of the two strategies is shown in the lower left panel, while the power demand of each component of the target controls strategy is shown in the lower right panel. Not surprisingly, the power demand is dominated by gas compression, especially during the first six months of the simulation.

5. Closed-loop workflow

The closed-loop workflow proposed in this paper is based on dividing the optimization period into smaller time intervals. We choose monthly intervals in this paper and introduce an interval index i (which goes from $i = 1$ to N_i , where N_i is the total number of months). The goal of the closed-loop workflow is to use a wind power forecast to reduce the CO₂ emissions during interval i by varying the well control rates from day to day, while minimizing deviations from the monthly target strategy (produced and injected volumes). We assume that in each interval i , a forecast of daily average wind speeds (which is used to compute available wind power) exists. For simplicity, uncertainty in the wind power forecast is neglected in this paper. To calculate a forecast of the daily average available wind power, we load the daily average wind speed data from the NORA10EI dataset ([Haakenstad et al., 2020](#)) from the period 2010 to 2015, and transform it into wind power using the power curve in [Fig. 7](#).

[Fig. 11](#) illustrates the steps in the proposed workflow. The first step in each interval is to calibrate the coarse model to data generated by simulating the fine-scale model with the target strategy for interval i . The calibration is performed with ensemble data assimilation using the ESM DA method ([Emerick and Reynolds, 2013](#)). The calibrated coarse model is then obtained from the mean of the posterior ensemble. Because permeability values span several orders of magnitude, the mean model is constructed by taking the exponential value of the log mean.

Before performing the short-term optimization based on the newly calibrated coarse model, the wind power forecast is loaded. In the short-term optimization, daily well controls are determined by minimizing the following objective function

$$J_{\text{short-term}}^i = \omega \left(\frac{M_{co2}^i - \hat{M}_{co2}^i}{\hat{M}_{co2}^i} \right) + (1 - \omega) \sqrt{\sum_{\alpha} \left(\frac{V_{\alpha}^i - \hat{V}_{\alpha}^i}{\hat{V}_{\alpha}^i} \right)^2} \quad (14)$$

where $\alpha \in \{o, g, wi, wp\}$ for oil, gas, water injection, and water production, respectively. All quantities with a hat indicate volumes or mass when simulating the target strategy with the coarse-scale model. For example, $(V_o^i - \hat{V}_o^i)/\hat{V}_o^i$ is the relative difference in the volume of produced oil in interval i . The parameter ω is a user-specified weight to prioritize the terms. The short-term objective is constructed such that $J_{\text{short-term}}^i = 0$ for the target strategy. The last term in $J_{\text{short-term}}^i$ penalizes any deviation (positive and negative) in the produced/injected volumes for month i , while the first term encourages a reduction in CO₂.

After the short-term optimization is performed, the fine-scale model is simulated with the short-term control strategy. This step is analogous to operating the reservoir field with the short-term controls. The cycle is repeated for each time interval, i , over the entire optimization period.

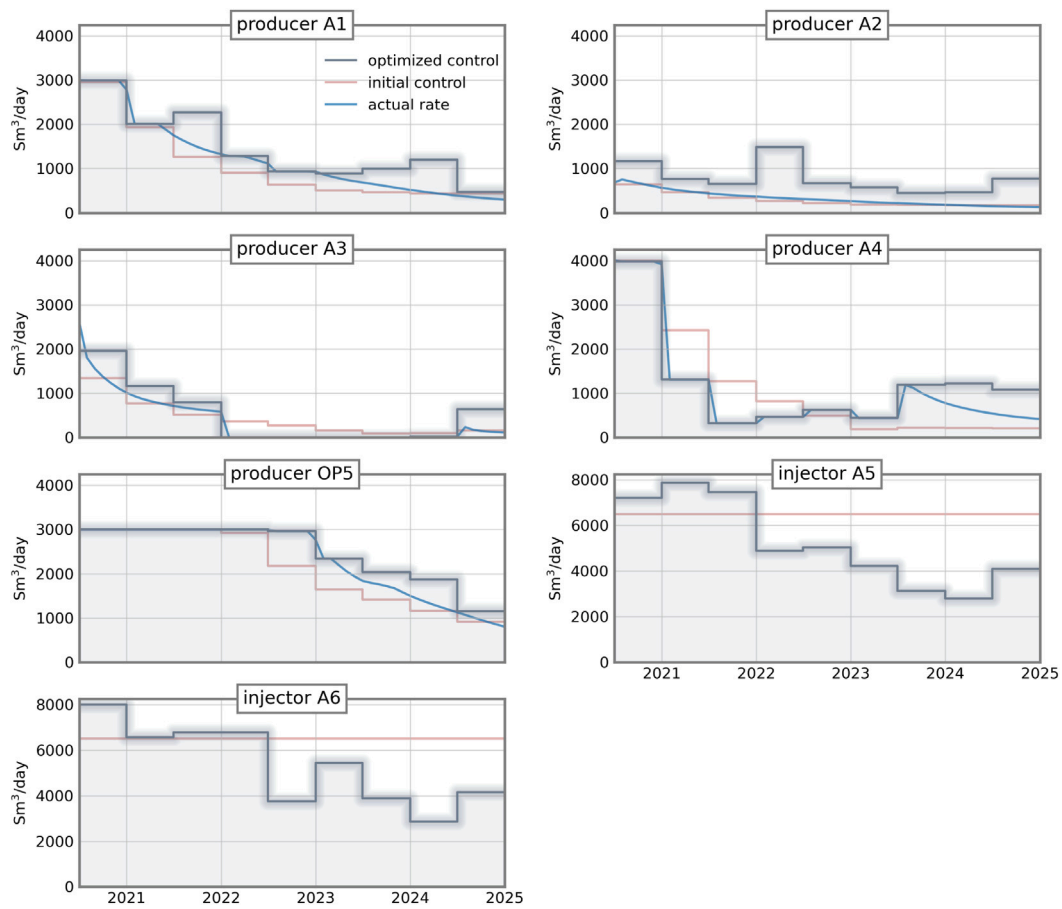


Fig. 9. Long-term target control strategy (gray lines) in each well of the Drogan reservoir. The red line in each panel is the initial control of the target optimization, and the blue lines are the actual oil production rates of the target control strategy (gray lines).

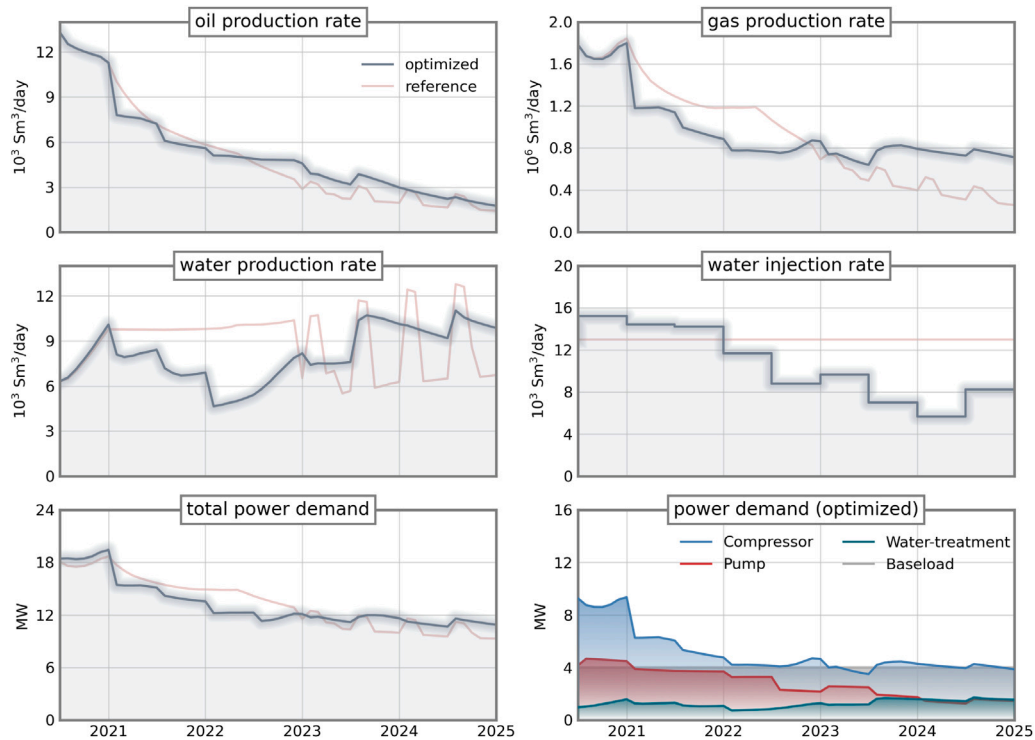


Fig. 10. The total oil, gas, and water production rates of the target strategy and the reference simulation. The total power demand of the two strategies is shown in the lower left panel, while the power demand of each component of the target controls strategy is shown in the lower right panel.

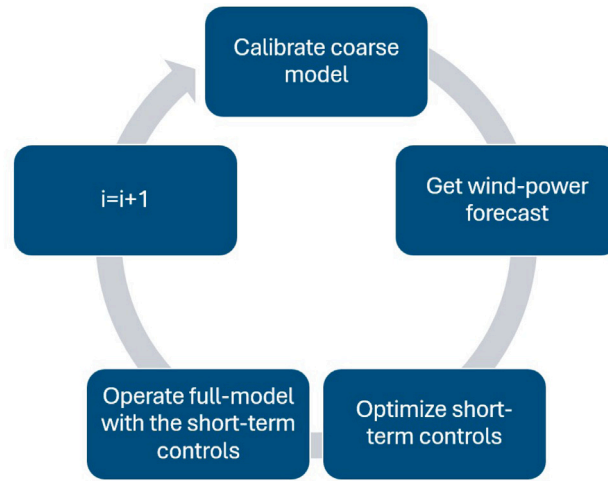



Fig. 11. The proposed closed-loop workflow.

Fig. 12. Scatter plot of the reduction in NPV and CO₂ of the workflow (with $\omega = 0.4, 0.5$, and 0.6) from the target strategy.

6. Numerical results

The workflow presented in the preceding section is run with three different values for the weight in Eq. (14): $\omega = 0.4, 0.5$, and 0.6 . For the calibration step of the workflow, the ESM DA method is used with the number of iterations set to five, and inflation values set to $1/5$. The prior ensemble is created by adding perturbations of a Gaussian field with mean equal to the posterior mean of the previous interval. For the first interval $i = 0$, the prior mean is given by the posterior mean from the history matching, see Section 2.2. Gaussian fields are sampled such that the standard deviation of the natural logarithm of the permeability (in x, y , and z direction) is 2, while the prior porosity is sampled with a standard deviation of 0.025. Both fields are sampled with spatial correlation, where the correlation length is two cells in the x and z direction, and five cells in the y direction.

The wind power forecast is created by loading daily average wind speeds from the NORA10EI dataset, then transforming the average wind speed to wind power using the power curve shown in Fig. 7. We multiply this by two, since we assume that two wind turbines are available.

For the short-term optimization, where the objective function $J_{\text{short-term}}^i$ is minimized, the ensemble-based optimization method as described in Fonseca et al. (2017) is employed. For ensemble-based optimization, samples are drawn from a Gaussian distribution where the user chooses the variance, and the mean is the current iterate of the optimization. A regression gradient of the objective function is calculated from the sample points, which is given by

$$\nabla J(x_k) \approx \frac{1}{N} \sum_{n=1}^N (J(x^n) - J(x_k))(x^n - x_k), \quad (15)$$

where J denotes the objective function, and the vectors $\{x^n\}_{n=1}^N$ represents the samples drawn from the Gaussian distribution. x_k denotes the current iterate of the optimization, where k is the iteration index. The gradient estimate from Eq. (15) is used in a gradient descent scheme, where the step-size is cut in half if a decrease in the objective function is not found. Upper and lower limits on the well rates are handled by sample truncation. The sample size for the gradient estimate is set to 200, and the standard deviation is set to 5% of the feasible range of the controls.

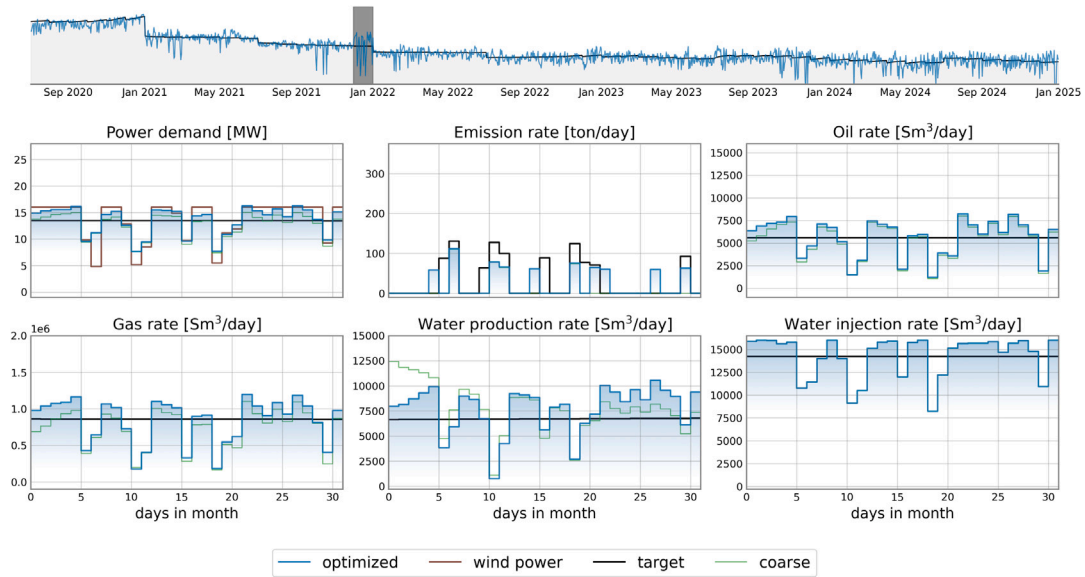


Fig. 13. Power demand, emission rate, and production data (FOPR, FGPR, FWPR, and FWIR) for the short-term strategy (blue) and target strategy (black) in interval $i = 18$ when $\omega = 0.5$. The green line shows what the coarse model predicts. The upper panel shows the timeline, and the shaded region indicates which interval the figure depicts.

The reduction (from the target) in CO_2 emissions and NPV of the three runs are depicted in Fig. 12. The run with $\omega = 0.6$ has an emission reduction slightly below 8%, and a decrease in NPV around 5%. Note that this emission reduction is in addition to the reduction that we get from including wind power, and running the target strategy without short-term optimization. With only gas turbines available (no wind power), the emissions resulting from the target strategy are 291 kilotonnes of CO_2 .

To illustrate the potential of the proposed workflow, Fig. 13 depicts the results from the interval $i = 18$ when $\omega = 0.5$. The upper panel shows the power demand for the entire 4.5-year period of the target (shown in black) and for the short-term strategies (shown in blue). The shaded region represents the interval we are currently looking at. The middle left panel shows the power demand of the short-term strategy (blue) and the target strategy (black) for the shaded region. The available wind power is also shown as the light red curve. The power demand that the coarse model predicts is shown in green. As the panel shows, the power demand of the optimized short-term operational strategy in this interval closely follows the available wind power. In contrast, the power demand of the target strategy is constant. This leads to significantly lower emissions in this interval. This can be seen in the center panel of the middle row, which depicts the emission rate of the different strategies. In this interval, the target strategy has an emission of 963 tonnes, while the emissions from the optimized short-term strategy are 699 tonnes of CO_2 , which is a 27% decrease from the target strategy.

The right panel of the middle row of Fig. 13 shows the oil production rate, while the lower panel shows the gas and water production and total water injection rates. The NPV in this interval is 65.3 million USD for the target strategy and 64.9 million USD from the optimized short-term strategy. This is only a 0.6% reduction in NPV (for this interval), which is significantly smaller than the 27% reduction in CO_2 emissions.

For the interval shown in Fig. 13, the workflow works as intended, reducing emissions significantly while not compromising the NPV, by taking advantage of the forecasted wind power. However, while there are several such intervals during the entire loop (for all three runs: $\omega = 0.4, 0.5$, and 0.6), many of the time intervals are not as successful as the one in Fig. 13. Interval $i = 37$ is an example of a less successful interval, it is shown in Fig. 14. In this interval, the CO_2 emissions and NPV of the

target strategy are 2208 tonnes and 37.7 million USD, respectively. In comparison, it is 2068 tonnes and 36.1 million USD for the optimized short-term strategy. This is a reduction in emissions of 6.3% and the reduction in NPV in this interval is 4.3%. The reason why the workflow is not as successful in this window can be the fact that the calibrated coarse model is not very accurate in this interval. This is evident by the difference of the gas rate given by the coarse model (represented by the green line in the lower left panel of Fig. 14) and the gas rate given by the fine-scale model (represented by the blue line). Since the gas rate predicted by the coarse model is almost half of the gas rate of the fine-scale model, the coarse model will underestimate the power demand, which causes the optimization to converge prematurely. Recall that gas compression is the most energy-demanding process modeled in this workflow.

7. Discussion and outlook

The workflow presented in this paper relies on several assumptions. A strong assumption is that we assume a perfect wind power forecast of the average available wind speed for a whole month. However, the workflow does not rely on the length of the intervals, i , and it is possible to choose shorter intervals, e.g., a week or 5 days. This choice makes the assumption more realistic, as justified by the forecast methodology developed in Shi et al. (2025), where they demonstrate the ability to produce a highly reliable forecast for wind power at every 15 min for up to three days ahead. Given that the workflow developed in this paper relies on daily averages, the assumption of a perfect forecast of daily average available wind power is not unrealistic.

Additionally, it is possible to include uncertainty in the wind power forecasts by generating multiple realizations of the forecast. Robust short-term optimization (see e.g., Nilsen et al., 2025) over the wind forecast ensemble can then be performed in each interval. In addition, we have only assumed one realization of the geology of the fine-scale reservoir model. In principle, one could include an ensemble of reservoir models and perform robust optimization over the ensemble of coarsened models calculated from the fine-scale models.

The proposed workflow is demonstrated on the Drogan test-reservoir model. Even though the short-term optimization shows limited improvements in many of the intervals of the closed loop, the results demonstrate the potential and value of a wind power forecast. In

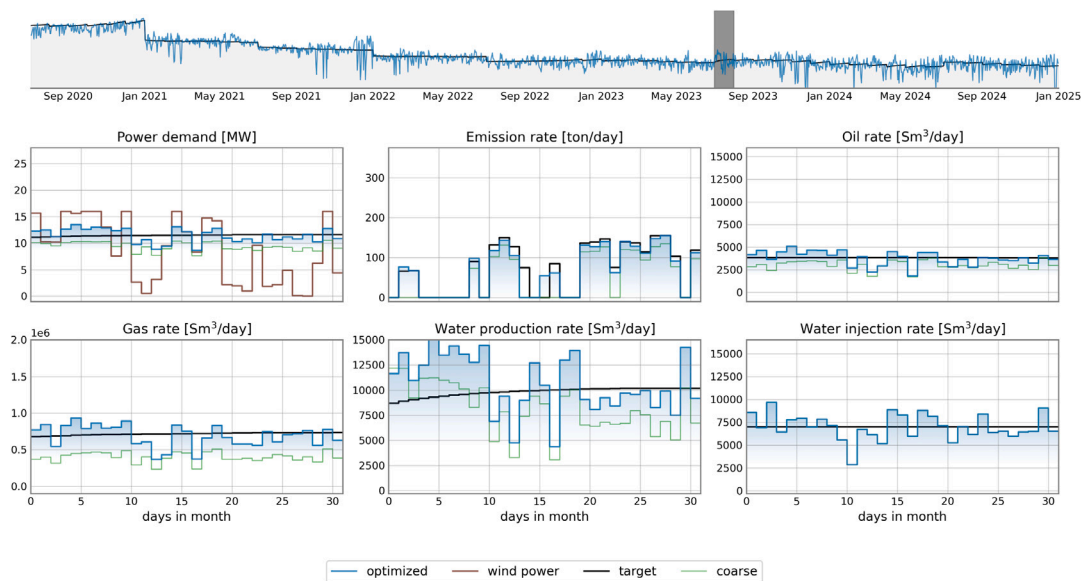


Fig. 14. Power demand, emission rate, and production data (FOPR, FGPR, FWPR, and FWIR) for the short-term strategy (blue) and target strategy (black) in interval $i = 37$ when $\omega = 0.5$. The green line shows what the coarse model predicts. The upper panel shows the timeline, and the shaded region indicates which interval the figure depicts.

one interval, the emissions were reduced by 27% while simultaneously only reducing the NPV 0.6%. It was also explained that the intervals where the workflow did not perform optimally were most likely due to a poor calibration of the coarse-scale model to the fine-scale model. For future work, we plan to test other proxy models, such as a simple regression model or the SubID method used in [van Essen et al. \(2013\)](#). This makes sense since the relationship between the controls and the objective function is not very nonlinear in a short time period (less than a month). We will also explore possible changes to improve the workflow in general, for example, by including constraints (such as setting a maximum deviation from the target) in the short-term optimization. Additionally, we could consider a storage option for excess wind power, to further reduce the dependency on gas turbines. Integration of wind power into energy systems with a storage option was considered in [Mohamad et al. \(2019\)](#) and [Teh and Cotton \(2016\)](#).

CRediT authorship contribution statement

Mathias M. Nilsen: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Formal analysis, Conceptualization. **Rolf J. Lorentzen:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Olwijn Leeuwenburgh:** Writing – review & editing, Supervision, Methodology, Conceptualization. **Andreas S. Stordal:** Writing – review & editing, Conceptualization. **Eduardo Barros:** Methodology, Conceptualization.

Declaration of competing interest

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

Acknowledgments

The authors acknowledge funding from the Centre of Sustainable Subsurface Resources (CSSR), grant nr. 331841, supported by the Research Council of Norway, Norway, research partners NORCE Norwegian Research Centre and the University of Bergen, and user partners Equinor ASA, Harbour Energy, Sumitomo Corporation, Earth Science Analytics, GCE Ocean Technology, and SLB Scandinavia.

Data availability

Data will be made available on request.

References

- Angga, I.G.A.G., Bellout, M., Kristoffersen, B.S., Bergmo, P.E.S., Slotte, P.A., Berg, C.F., 2022. Effect of CO₂ tax on energy use in oil production: waterflooding optimization under different emission costs. *SN Appl. Sci.* 4 (313), <http://dx.doi.org/10.1007/s42452-022-05197-4>.
- Emerick, A.A., Reynolds, A.C., 2013. Ensemble smoother with multiple data assimilation. *Comput. Geosci.* 55, 3–15. <http://dx.doi.org/10.1016/j.cageo.2012.03.011>.
- Equinor, 2025a. Drogon conceptual description. <https://webviz-subsurface-example.azurewebsites.net/drogon-conceptual-description>.
- Equinor, 2025b. Information for analysts. <https://www.equinor.com/investors/information-for-analysts#downloads>. (Accessed: 20 March 2025).
- Evensen, G., Vossepoel, F.C., van Leeuwen, P.J., 2022. *Data Assimilation Fundamentals: A Unified Formulation of the State and Parameter Estimation Problem*. Springer Nature, Cham.
- Fonseca, R.R.-M., Chen, B., Jansen, J.D., Reynolds, A., 2017. A Stochastic Simplex Approximate Gradient (StoSAG) for optimization under uncertainty. *Internat. J. Numer. Methods Engrg.* 109 (13), 1756–1776. <http://dx.doi.org/10.1002/nme.5342>.
- Fonseca, R.M., Rossa, E.D., Emerick, A.A., Hanea, R.G., Jansen, J.D., 2020. Introduction to the special issue: Overview of OLYMPUS Optimization Benchmark Challenge. *Comput. Geosci.* 24, 1933–1941. <http://dx.doi.org/10.1007/s10596-020-10003-4>.
- Guo, B., Lyons, W.C., Ghalambor, A., 2007. *Transportation Systems*. In: Guo, B., Lyons, W.C., Ghalambor, A. (Eds.), *Petroleum Production Engineering*. Gulf Professional Publishing, Burlington, pp. 133–158. <http://dx.doi.org/10.1016/B978-075068270-1/50016-5>.
- Haakenstad, H., Breivik, Ø., Reistad, M., Aarnes, O.J., 2020. NORA10EI: A revised regional atmosphere-wave hindcast for the North Sea, the Norwegian Sea and the Barents Sea. *Int. J. Climatol.* 40 (10), 4347–4373. <http://dx.doi.org/10.1002/joc.6458>.
- Huo, Q., Li, X., Deng, H., Wang, W., Xu, F., Lu, Y., 2023. Research on the optimal allocation of new energy connected to electrolytic aluminum. In: 2023 IEEE 6th International Electrical and Energy Conference. CIEEC, pp. 2521–2526. <http://dx.doi.org/10.1109/CIEEC58067.2023.10166475>.
- Landa-Marbán, D., 2024. Csr-tools/pycopm: 2024.04 (v2024.04). <http://dx.doi.org/10.5281/zenodo.12740839>, Zenodo.
- Liu, Y., Zhu, Y., Bi, X., Liu, X., Zhang, Y., Wang, Y., 2025. Physics-informed neural networks for PDE problems: a comprehensive review. *Artif. Intell. Rev.* <http://dx.doi.org/10.1007/s10462-025-11322-7>.
- Mohamad, F., Teh, J., Abunima, H., 2019. Multi-objective optimization of solar/Wind penetration in power generation systems. *IEEE Access* 7, 169094–169106. <http://dx.doi.org/10.1109/ACCESS.2019.2955112>.

- Musial, W., Beiter, P., Smith, A., Tegen, S., 2016. Potential Offshore Wind Energy Areas in California: An Assessment of Locations, Technology, and Costs. Technical Report NREL/TP-5000-67414, National Renewable Energy Laboratory, <http://dx.doi.org/10.2172/1338174>.
- My Engineering Tools, 2025. Compressor power requirement. https://myengineeringtools.com/Compressors/Tools_Compressor_Power.html. (Accessed: 12 March 2025).
- Nilsen, M.M., Lorentzen, R.J., Stordal, A.S., Leeuwenburgh, O., Barros, E., 2025. Wind-powered reservoir management with application to robust multi-objective optimization. *Comput. Geosci.* 29, <http://dx.doi.org/10.1007/s10596-025-10370-w>.
- Norwegian Petroleum, 2025. Energy calculator. <https://www.norskipetroleum.no/en/calculator/about-energy-calculator/>. (Accessed: 12 February 2025).
- Offshore Norge, 2024. Environmental Report 2023. Technical Report, Offshore Norge, <https://info.offshorenorge.no/environmental-report-2023>.
- Oliver, D.S., Raanes, P.N., Skorstad, A., Sætrom, J., 2025. Multi-objective reservoir production optimization: Minimizing CO2 emissions and maximizing profitability. *Geoenergy Sci. Eng.* 244, 213396. <http://dx.doi.org/10.1016/j.geoen.2024.213396>.
- Oliver, D., Sætrom, J., Skorstad, A., Raanes, P., Saksvik, T., Kolbjørnsen, O., 2024. Robust multi-objective production optimization with CO2 emissions reduction. In: European Conference on the Mathematics of Geological Reservoirs. <http://dx.doi.org/10.3997/2214-4609.202437098>.
- Rasmussen, P., Kurz, R., 2009. Centrifugal compressor applications - Upstream and midstream. In: Thirty-Eighth Turbomachinery Symposium. <http://dx.doi.org/10.21423/R1664Q>.
- Sandve, T., Lorentzen, R., Marban, D., Fossum, K., 2024. Closed-loop reservoir management using fast data-calibrated coarse models. In: European Conference on the Mathematics of Geological Reservoirs. <http://dx.doi.org/10.3997/2214-4609.202437071>.
- Shi, J., Teh, J., Lai, C.-M., 2025. Wind power prediction based on improved self-attention mechanism combined with Bi-directional Temporal Convolutional Network. *Energy* 322, 135666. <http://dx.doi.org/10.1016/j.energy.2025.135666>.
- Teh, J., Cotton, I., 2016. Reliability impact of dynamic thermal rating system in wind power integrated network. *IEEE Trans. Reliab.* 65 (2), 1081–1089. <http://dx.doi.org/10.1109/TR.2015.2495173>.
- The Engineering ToolBox, 2025. Combustion of fuels - Carbon dioxide emission. https://www.engineeringtoolbox.com/co2-emission-fuels-d_1085.html. (Accessed: 12 February 2025).
- van Essen, G.M., Van den Hof, P.M., Jansen, J.-D.D., 2013. A two-level strategy to realize life-cycle production optimization in an operational setting. *SPE J.* 18 (06), 1057–1066. <http://dx.doi.org/10.2118/149736-PA>.