

ORIGINAL PAPER

Wind-powered reservoir management with application to robust multi-objective optimization

Mathias M. Nilsen^{1,2} · Rolf J. Lorentzen¹ · Andreas S. Stordal¹ · Olwijn Leeuwenburgh³ · Eduardo Barros³

Received: 29 November 2024 / Accepted: 26 June 2025
© The Author(s) 2025

Abstract

This paper addresses the challenge of incorporating offshore wind power into reservoir management. Traditionally, oil and gas production is powered by gas turbines. While stable, gas turbines are a major source of CO₂ emissions. In contrast, wind power produces power with minimal emissions. However, due to its high variability and uncertainty, including it in the optimization of operational strategies over extended periods can be challenging. In this paper, the optimization of production strategies over an ensemble of realistic wind power series is investigated. The ensemble is generated by a mathematical model consisting of an autoregressive model with a seasonal trend. The model is conditioned on relevant wind speed data from the North Sea with Bayesian inference. The wind speed data is selected from the open-access NORA10EI dataset. The methodology developed in this paper is applied to a multi-objective optimization problem, focusing on studying the tradeoff between profit and emissions. A benchmark test reservoir model and a detailed CO₂ emissions calculator are employed. In this scenario, wind power is combined with traditional gas power, and all results are compared with a reference where only gas power is used. The experiment indicates that it is not possible to reduce emissions by 40% without the use of wind power.

Keywords Reservoir management · Wind power · Multi-objective optimization · Ensemble optimization

1 Introduction


In 2022, the Norwegian oil and gas industry emitted around 11.2 million tonnes of CO₂ equivalents [1]. According to the Norwegian Offshore Directorate, this accounted for around a quarter of Norway's total greenhouse gas emissions that year. Of those eleven million tonnes, approximately 84% (see Fig. 1) was estimated to come from gas turbines, generating electricity for various platform facilities and powering the oil and gas production. Over the recent years, much effort has been put into reducing the use of gas turbines by introducing offshore wind. An example is the Hywind Tampen wind farm, with its eleven 8MW wind turbines supplying power to the Gullfaks and Snorre fields [2]. Hywind Tampen covers about 35% of the power demand of the two fields, and it is projected to reduce CO₂ emissions by 200 kilotonnes per year [3].

Besides reducing emissions, the petroleum sector is incentivized economically to include renewables through carbon emission taxes. Including the cost of emissions trading, the total cost associated with CO₂ emissions for the Norwegian oil and gas industry is projected to reach 2000 NOK per tonne of CO₂ in 2030 [1]. Angga et al. [4] studied the effect of CO₂ tax on water flooding optimization and found that increasing tax affects the emissions more than the profits from production by encouraging more energy-efficient reservoir drainage strategies.

Rolf J. Lorentzen, Andreas S. Stordal, Olwijn Leeuwenburgh, and Eduardo Barros contributed equally to this work.



Fig. 1 The relative size of CO₂ emitters from the Norwegian continental shelf in 2022 [1]



In recent decades, numerous studies have focused on optimizing the economic performance of reservoirs, as seen in the works of [5–7]. More recently, research has started to explore the trade-off between profits and emissions resulting from production. As mentioned in the previous paragraph, [4] investigated the impact of introducing a CO₂ tax. Additionally, [8] and [9] applied the weighted sum method from multi-objective optimization, demonstrating that it is possible to achieve a significant reduction in emissions without notably compromising profits. However, how the operational strategy changes the trade-off when wind power is included in the optimization is not clear, which is the focus of this study.

From a reservoir management and optimization perspective, the biggest challenge with wind power is its high uncertainty and variability, making it hard to optimize production strategies over longer periods of time. This paper addresses this problem by generating an ensemble of realistic wind speed (wind power) profiles that can be included when optimizing various control strategies. The wind speed ensemble is generated by a mathematical model, where the model parameters are conditioned on data from the North Sea with Bayesian inference. An empirical power curve model of an 8MW offshore wind turbine is used to derive power profiles from the wind speed profiles. In the optimization, traditional gas turbines cover the residual power demand not covered by wind power.

The methodology developed in this paper is applied to a bi-objective optimization case, where a realistic benchmark test reservoir is used. The weighted sum method for multi-objective optimization is utilized to find a set of Pareto solutions. In this scenario, the two objective functions are Net Present Value (NPV), which is the economic putout of the production, and CO₂ emissions. A stochastic ensemble-based gradient method well suited for robust optimization problems is employed. The resulting Pareto front is compared to a second Pareto front optimized without wind power (with only power from gas turbines available) for reference.

The outline of the paper is as follows. In Section 2, the wind speed data and the mathematical model used for generating the wind power ensemble are discussed. Section 3 describes the ensemble-based gradient and how it is used for robust optimization. A brief introduction to multi-objective optimization and the weighted sum method is also covered here. The Drogon reservoir is introduced in Section 4, and the simulation workflow is explained. In Section 5, the methodology is applied to the scenario of minimizing emissions while maximizing NPV. Finally, the paper is summarized in Section 6, where some final remarks are made.

2 Wind power simulation

The top panel of Fig. 2 shows daily wind speeds from August 1, 2010 to January 1, 2015 at a height of 80 meters. The location is at a latitude of 61.3°N and longitude of 2.3°E, approximately where the Hywind Tampen wind farm is located. The data is from the open-access dataset NORA10EI [10], which is an atmospheric hindcast for the Norwegian, North, and Barents Sea. NORA10EI contains data from 1979 to 2017 with a model resolution of 0.1°.

Among many sources, Brown et al. [11] discuss that the stationary distribution of a wind speed time series is described well by a Weibull distribution. This is also supported by the Weibull fit of the data shown in the top right corner of Fig. 2. Motivated by [11], the approach taken in this paper is to transform the stationary distribution of

Fig. 2 Wind speed from the NORA10EI dataset (upper panel) with a histogram of the data together with a Weibull fit. The data is transformed to be Gaussian distributed (lower panel)



the data to a standard Gaussian, $\mathcal{N}(0, 1)$. This is done by transforming each wind speed value with the cumulative distribution function (CDF) of the Weibull distribution. Then, the standard Gaussian inverse CDF transform is applied, resulting in data with a standard Gaussian stationary distribution. The lower panel of Fig. 2 shows the transformed Gaussian data with the stationary distribution to the right.

An autoregressive model is then fitted to the Gaussian data and then used to simulate Gaussian wind speed profiles. The simulated time series is finally transformed back to the original Weibull distribution from Fig. 2. This ensures that all simulated wind speed series have the same stationary distribution as the data and that all generated wind speed values are positive. This approach differs from the one taken in [11], where they did not directly transform the data to a Gaussian distribution. Instead, they manipulate the parameters of the Weibull distribution to mimic a Gaussian shape.

The data in Fig. 2 exhibits clear seasonal variations, with generally stronger wind speeds during the winter months. Because of this, the autoregressive model proposed in this paper includes a cosine term to model seasonal variability. The model is given by

$$x_{t+1} = \rho_0 + \rho_1 x_t + \varepsilon_t + a \cos\left(\frac{2\pi}{365\lambda}t + \phi\right), \quad (1)$$

where $\varepsilon_t \sim \mathcal{N}(0, \sigma_m^2)$. Bayesian inference is then applied to the model parameters,

$$\theta = [\rho_0, \rho_1, a, \lambda, \phi, \sigma_m]^\top, \quad (2)$$

to get a posterior distribution conditioned on the data. The posterior distribution, $p(\theta|d)$, is given by Bayes' theorem

$$p(\theta|d) = \frac{p(d|\theta)p(\theta)}{p(d)}, \quad (3)$$

where $d \in \mathbb{R}^k$ is a vector containing the k Gaussian data values, $p(d|\theta)$ is the likelihood function and $p(\theta)$ is the prior distribution of θ . The likelihood is assumed to be Gaussian, such that

$$p(d|\theta) = \frac{1}{\sqrt{(2\pi)^k \det(\Sigma)}} \exp\left[-(d - x(\theta))^\top \Sigma^{-1}(d - x(\theta))\right], \quad (4)$$

where $x(\theta) \in \mathbb{R}^k$ is a vector containing the simulated time series, and the covariance matrix is given by $\Sigma = \sigma_m^2 I$. The chosen priors for the parameters θ in Eq. 2 are listed in Table 1. The priors were selected based on informed

Table 1 The marginal prior distribution for each parameter in the model in Eq. 1

	ρ_0	ρ_1	a	λ	ϕ	σ_m
prior	$\mathcal{N}(0, 0.1^2)$	$\mathcal{N}(0.9, 0.1^2)$	$\mathcal{N}(1, 0.1^2)$	$\mathcal{N}(1, 0.1^2)$	$\mathcal{N}(0, \frac{\pi}{2}^2)$	Half-Norm($\sigma = 2$)

estimates, with sufficiently large variances to ensure that the posterior distributions would be primarily shaped by the observed data.

The Bayesian inference is performed with the python package PyMC [12], following the example in [13]. The marginal prior and posterior distributions of the model parameters are shown in Fig. 3. Because the posterior is sampled with a Markov Chain Monte Carlo (MCMC) method in PyMC, the marginal posteriors in Fig. 3 are represented by KDEs (Kernel Density Estimates).

A random sample of the model parameters is now drawn from the posterior distribution. Each sample point is then used to simulate a time series of Gaussian-distributed wind speeds before being transformed back to a Weibull distribution, as described earlier. Figure 4 shows six examples of generated wind speed profiles. Each wind profile is transformed to a power series with the power curve from an 8MW offshore wind turbine, found in [14] (data for the power curve are available at [GitHub](#)). The power curve is depicted in Fig. 5. In total, two wind power ensembles of size 50 are generated. One is to be used in the optimization, and the other is used to evaluate the optimized results.

2.1 Remark

In this paper, directional changes in wind are not considered. Wind turbines are installed with *yaw systems*, responsible for turning the turbines when the wind direction changes. A typical angular yaw speed is 0.3° per second [15]. Defining the angle between the wind velocity and the turbine area as γ , the power output can be approximated by [16]

$$P(\gamma) = P_0 \cos^3 \gamma, \quad (5)$$

where $P_0 = 8$ MW is the power output when $\gamma = 0$. Assuming that the turbine needs to turn from $\gamma_1 = 45^\circ$ to $\gamma_2 = 0^\circ$, with the angular speed of $\omega_d = 0.3^\circ$ per second, the total time it takes to turn is $t_d = 150$ seconds.

Fig. 3 Marginal prior (red) and posterior (blue) distribution of the model parameters θ





Fig. 4 Six generated windspeed profiles from the model in Eq. 1



During this time, the energy loss, ΔE , is given by

$$\Delta E = \int_{t=0}^{t=t_d} P_0 \cos^3 \gamma(t) \cdot dt = \int_{t=0}^{t=t_d} P_0 \cos^3 \left(\frac{\pi}{180} \omega_d \cdot t \right) \cdot dt \approx 0.25 \text{ MWh.} \quad (6)$$


Assuming this occurrence happens once every hour, we find that the total energy loss per day (over 24 hours) is 6 MWh. In an ideal scenario, the total energy output would be 8 MW multiplied by 24 hours, which equals 192 MWh. Therefore, the energy lost due to changing wind direction accounts for only 3.125%. It is possible that wind direction changes more frequently during the winter months than in the summer. This could lead to higher energy losses in winter, which might affect the shape of the power profiles. However, since this does not alter the workflow presented in this paper, it is not included in the analysis.

3 Robust ensemble optimization

Ensemble-based optimization (EnOpt) is one of the popular gradient methods for optimizing well controls in reservoir management [6, 7]. Denoting the objective function as $f(u)$, where u is the control vector to be determined by optimization, the ensemble-based gradient is given by

$$\Sigma \nabla f(u) \approx \frac{1}{N} \sum_{n=1}^N (f(u^n) - f(u))(u^n - u), \quad (7)$$

Fig. 5 Power curve for a 8MW wind turbine in [14]



where the vectors $\{u^n\}_{n=1}^N$ are samples from a Gaussian distribution with mean u and a covariance matrix Σ . The pre-multiplication of Σ in Eq. 7 is viewed as pre-conditioning of the gradient. This term can be removed by multiplying the equation with Σ^{-1} if not desirable. Once the gradient is estimated, the control vector is updated iteratively by the gradient descent method,

$$u_{j+1} = u_j - \alpha_j g_j, \quad (8)$$

where g_j is the gradient estimate in Eq. 7 at the current iteration j , and α_j is the step size.

There are several reasons to use EnOpt in optimization for reservoir management. Firstly, it is nonintrusive, meaning the inner workings of the objective function are irrelevant when estimating the gradient. This quality is attractive when a dynamic simulation, such as a reservoir flow simulation, is needed to evaluate the function. Secondly, EnOpt can deal with robust optimization problems. Robust optimization is when the objective function depends on two inputs: the control to be determined, u , and an uncertain variable, χ . In reservoir management, this is often some geological property like permeability or porosity. Let χ^n , where n goes from 1 to N , be different realizations of the variable χ . In the gradient estimate of EnOpt, each χ^n is then paired with a control perturbation u^n such that the estimate in Eq. 7 becomes

$$\Sigma \nabla f(u) \approx \frac{1}{N} \sum_{n=1}^N \left(f(u^n, \chi^n) - f(u, \chi^n) \right) (u^n - u). \quad (9)$$

In this manner, EnOpt can efficiently estimate a gradient with the uncertainty in χ included. At each iteration of the optimization, the proposed control u_{j+1} is evaluated on the average objective function $\overline{f(u_{j+1})} = N^{-1} \sum_{n=1}^N f(u_{j+1}, \chi^n)$, and a convergence check is performed.

3.1 Multi-objective optimization

The following is a brief summary of the weighted sum method for solving multi-objective optimization presented by Goodarzi et al. in [17]. In multi-objective optimization, the goal is to minimize several objective functions simultaneously. However, if the objectives compete with each other, it is generally not possible to find a single optimum value. Therefore, one tries to find a set of solutions called a *Pareto front* or a *Pareto optimal set*. The solutions in the Pareto front are optimal in the sense that decreasing the value of one objective function leads to a simultaneous increase in at least one of the other objectives.

Given several objective functions, $f_l(u)$, where $l = 1, \dots, L$, the weighted sum method of multi-objective optimization combines the objectives to form a single objective function

$$F(u) = \sum_{l=1}^L \omega_l f_l(u), \quad (10)$$

where $\omega_l \in [0, 1]$ are weights assigned to the different objectives, such that $\sum_{l=1}^L \omega_l = 1$. In this paper, we only consider two objective functions, making the optimization bi-objective, and the combined function in Eq. 10 becomes


$$F(u) = \omega f_1(u) + (1 - \omega) f_2(u). \quad (11)$$

The Pareto front is then found by optimizing $F(u)$ for a range of ω values.


4 The Drogon field and simulation workflow

Drogon is a benchmark test reservoir developed by Equinor [18] for testing optimization methods. The field measures approximately $3.6\text{km} \times 4.5\text{km} \times 0.1\text{km}$ [19] and is divided into a $46 \times 73 \times 31$ grid with 71 575 active cells. The simulation of Drogon is divided into two periods: one period for history matching followed by a prediction/optimization period starting July 1, 2020, and lasting until January 1, 2025. The topology of the Drogon reservoir is depicted in Fig. 6a, while Fig. 6b shows the porosity, and Fig. 6c shows the permeability in the x-direction. Drogon has two injection wells labeled A5 and A6, and four production wells labeled A1 to A4. In addition, an extra production well, denoted OP5, is added at the beginning of the optimization period. This well is part of the schedule defined by Equinor, and its purpose is to see what effect an additional well can have on the predictions of a history-matched ensemble. However, this is not studied here, but the well is still included. An ensemble of history-matched porosity and permeability realizations is provided by Equinor for the Drogon field, and in this paper, fifty of them are used. Figure 6b and 6c depict one of these realizations.


The controls to be determined by optimization in the next section are the well water injection rate (WWIR) in the two injectors A5 and A6, and the target field oil production rate (FOPR), both adjusted monthly during the prediction period. This results in $3 \times 54 = 162$ (three control types times 54 months) controls to be determined. When it comes to the upper bounds of the controls, the maximum injection rate for each well is $8000 \text{ Sm}^3/\text{day}$. Furthermore, the producers A2, A3, and A4 have a maximum oil production rate of $4000 \text{ Sm}^3/\text{day}$, and A1 and OP5 have a limit of $3000 \text{ Sm}^3/\text{day}$. This means that the upper bound for the target oil production rate for the field is $18000 \text{ Sm}^3/\text{day}$.



(a) Topography of Drogon




(b) A single realization of the porosity



(c) A single realization of the permeability in x-direction

Fig. 6 The Drogon reservoir

Fig. 7 Flow diagram explaining the simulation workflow



To estimate the emissions from Drogon's operation, a simple model of the top-side facility of Drogon is chosen, where three components of power consumption are present. The three components are a constant baseload power demand of 4MW, a water injection pump for the two injectors, and a gas compressor for the produced gas. The power demand of the pump and the compressor varies with the rate and the difference in inlet and outlet pressure. The emissions calculator *eCalc* [20] is used to model the power demand and CO₂ emissions. To cover the power demand of the production, it is assumed that one 8MW wind turbine is available and that gas turbines are available to cover the remaining load when needed.

The simulation workflow is illustrated in Fig. 7, where the first input u , is a vector containing the monthly WWIR for A5 and A6, and target FOPR. The second input is the permeability in the x,y, and z directions (denoted PERMX, PERMY, and PERMZ), and the porosity, denoted PORO. The third input is P_{wind} , which is a vector containing daily available wind power. The OPM simulator [21] is used to simulate the reservoir flow, and it outputs the volume of produced oil, gas, and water (denoted $V_{\text{oil-prod}}$, $V_{\text{gas-prod}}$, and $V_{\text{wat-prod}}$) each month. *eCalc* computes the total power demand, $P_{\text{tot-dem}}$ of the pump, compressor, and baseload. The power load required by the gas turbines, denoted $P_{\text{gas-dem}}$, is calculated by subtracting the available wind power from $P_{\text{tot-dem}}$. $P_{\text{gas-dem}}$ is then used to determine the gas fuel consumed by the gas turbines, denoted $V_{\text{gas-dem}}$, which is subtracted from the volume of produced gas from the reservoir. Lastly, the monthly emissions, M_{co_2} are determined (with *eCalc*) from $V_{\text{gas-dem}}$.

5 Numerical application: emissions vs. NPV

The workflow developed in the preceding sections is now applied to a bi-objective optimization case where the two objective functions are CO₂ emissions and Net Present Value (NPV). Uncertainty in both geology and wind power is considered in such a manner that each wind power realization, P_{wind}^n , is paired with one realization of porosity, denoted Φ^n , and permeability K^n . This means that χ^n in Eq. 9 is defined as $\chi^n = [P_{\text{wind}}^n, \Phi^n, K^n]$. The number of geological and wind power realizations is 50, which is also the sample size for the ensemble gradient described in Section 3. Higher ensemble sizes generally improve gradient quality [22]. However, a preliminary test on the Drogon model shows that a size of fifty gives satisfactory results.

NPV is defined by the revenue from oil and gas sales minus the cost of water disposal and injection. By defining c_o and c_g as the price of oil and gas per standard cubic meter (Sm³), and letting c_{wp} and c_{wi} be the cost of water production and injection per Sm³, and c_{co_2} the tax per kg of emitted CO₂, the NPV function, denoted f_{NPV} , can be written as

$$f_{\text{NPV}} = \sum_{i=1}^{N_c} \frac{c_o V_{\text{oil-prod}}^i + c_g V_{\text{gas-prod}}^i - c_{\text{wp}} V_{\text{wat-prod}}^i - c_{\text{wi}} V_{\text{wat-inj}}^i - c_{\text{co}_2} M_{\text{co}_2}^i}{(1+r)^{t_i/365}}, \quad (12)$$

where N_c is the number of intervals the controls are changed (number of months in this scenario), and r is the discount rate chosen to 8% per year. t_i is the number of days passed since $i = 0$. The economic parameters of the

Table 2 Specific values of the economic parameters in the NPV

c_o	c_g	c_{wp}	c_{wi}	c_{co_2}
600 \$/Sm ³	17 \$/Sm ³	38 \$/Sm ³	18 \$/Sm ³	0.15 \$/kg

NPV are given in Table 2. Additionally, the total mass of emitted CO₂ during the simulation period can be written as

$$f_{CO_2} = \sum_{i=1}^{N_c} M_{co_2}^i. \quad (13)$$

5.1 Reference simulation

As a reference, the workflow in Fig. 7 is simulated over the entire ensemble (χ^n) with a constant injection rate of 6000 Sm³/day in both injectors and maximum target FOPR of 18000 Sm³/day, meaning that the producers are allowed to produce as much as possible. The resulting oil, gas, and water production rate and CO₂ emission rate are depicted in Fig. 8. The shaded region shows the variability over the ensemble (porosity, permeability, wind power), while the white line shows the ensemble mean. Notice how the CO₂ emission rate oscillates with the seasons, resulting from the seasonal variability of the wind speed profiles. The scenario depicted in Fig. 8 results in an average NPV of \$25.7 billion and a total (average) of 246.3 kilotonnes of CO₂ emissions (denoted f_{NPV}^{ref} and $f_{CO_2}^{ref}$).


5.2 Optimization

For the optimization, the initial control, u_0 , is chosen to be a constant injection rate of 6000 Sm³/day in both injectors and the target FOPR set to the resulting oil production rate from the reference simulation (Fig. 8). The weighted-sum objective (Eq. 11) is written as

$$F(u, \chi^n) = \omega \frac{f_{CO_2}(u, \chi^n)}{f_{CO_2}^{ref}} + (1 - \omega) \frac{-f_{NPV}(u, \chi^n)}{f_{NPV}^{ref}}, \quad (14)$$

such that the two terms are dimensionless and comparable in size. The minus sign in front of f_{NPV} is there because the optimization algorithm is formulated in terms of minimization.

Fig. 8 Upper left: The field oil production rate. Upper right: Field gas production rate. Lower left: Field water production rate. Lower right: CO₂ emission rate



The gradient descent method (Eq. 8) is employed with an ensemble-based gradient as described in Section 3. At each iteration of the optimization, the Line Search method (described in Chapter 3 of [23]) is used to set the step size, and the initial step size is set to $\alpha_j^{\max} = 0.25 / \|g_j\|_{\infty}$, where g_j is the ensemble gradient estimate at iteration j . The variance of the sampling for the ensemble gradient is chosen such that the standard deviation is 5% of the feasible interval for each control. The maximum number of iterations is 30, and the maximum number of inner iterations (number of step size cuts) for the Line Search method is set to 5. One resampling of the ensemble gradient is allowed on any given iteration if the Line Search method does not find an improvement in the objective function. Lastly, the controls are sampled with a time correlation such that the correlation of two controls u_{t_k} and u_{t_l} are given by

$$\text{Corr}(u_{t_k}, u_{t_l}) = \rho^{|k-l|}, \quad (15)$$

where $\rho = 0.75$ is the correlation of two consecutive time intervals.

The main panel of Fig. 9 depicts two sets of solutions of the weighted sum objective function, each with five different values of the weight, ω . The set of solutions to the left of the main panel is determined with the wind power ensemble as described in Fig. 7, and the set to the right is determined with only gas power available. A panel that zooms in on the region of interest for each set is also shown (upper panel for the solutions with wind and gas power combined, and to the left for the solutions with only gas power.). For each solution of a particular value of ω , the small dots represent different ensemble members, χ^n , while the big dots represent the ensemble mean. The same data is given in Table 3, which gives the ensemble mean $\pm$ one standard deviation. The table also gives the emission intensity, I_{CO_2} , of each solution. Emission intensity is defined as the ratio between the amount of CO₂ emitted and the total volume of oil and gas produced. I_{CO_2} is given in units of kilogram per tonne of oil equivalents (kg/toe) in this paper. The resulting controls for each solution are depicted in Fig. 10, where each row shows the controls for a particular value of ω . The first and middle columns depict the injection rate in A5 and A6, respectively, and the last column depicts the target group production rate. The colored line in each panel shows the rate for the solutions with wind power, while the grey line is for the case with no wind power (only gas).

Fig. 9 Two sets of Pareto solutions of the weighted sum objective in Eq. 14, with $\omega = 0.0, 0.25, 0.5, 0.75$ and 1.0 . The solution set to the left is solved over the wind power combined with gas power, while the set to the right is solved with only gas power available. The dashed lines are 3-degree polynomial fits. CO₂ emissions are given on the x-axis, while the NPV is on the y-axis. The upper and right panels zoom in on regions of interest in the two sets




Table 3 The numerical values depicted in Fig. 9

	$\omega = 0.0$	$\omega = 0.25$	$\omega = 0.5$	$\omega = 0.75$	$\omega = 1.0$	
NPV [Billion USD]	27.27 ± 0.73	26.79 ± 0.73	26.04 ± 0.56	25.59 ± 0.053	1.82 ± 0.02	Wind + Gas
	26.66 ± 0.72	26.5 ± 0.67	26.35 ± 0.65	24.46 ± 0.52	0.07 ± 0.0	Gas
CO_2 [kilo tonnes]	223.71 ± 4.87	212.03 ± 4.65	200.61 ± 5.45	194.55 ± 5.90	113.83 ± 4.33	Wind + Gas
	339.94 ± 2.53	329.39 ± 2.81	326.22 ± 2.82	308.83 ± 2.97	197.91 ± 0.06	Gas
I_{CO_2} [kg/toe]	28.84 ± 0.7	28.69 ± 0.85	27.11 ± 0.77	27.15 ± 0.84	151.06 ± 5.77	Wind + Gas
	43.25 ± 0.72	43.95 ± 0.73	43.84 ± 0.70	44.12 ± 0.81	1943.49 ± 1.07	Gas

The table also shows the emission intensity of each solution




Fig. 10 The determined controls of the optimization. Each row represents a distinct value of the weight ω . Each column is a different type of control; the left and middle columns are the injection rates in A5 and A6, and the column to the right is the group control rate. The colored line on each panel is the control for the case where wind and gas power are combined, while the gray line is the control for when only gas power is available. Time is on the x-axis, and the rate given in Sm^3/day is on the y-axis

5.3 Discussion

The first thing to notice in Fig. 9 is the difference in emissions values (on the x-axis) of the two solution sets. By just introducing wind power, the total emissions are drastically reduced. In this case, the reduction is on the order of 100 kilotonnes. However, other qualitative differences can also be observed in the two sets. The trade-off between NPV and emissions looks more linear in the set with wind power than in the one with only gas power when ω is reduced from 0.0 to 0.75, which is positive as it means that NPV will not drop off as quickly as emissions are reduced. Additionally, the resulting controls with available wind power seem to have clearer seasonal variability (see Fig. 10) than those without.

When the objective function only contains the emissions ($\omega = 1.0$), the solution with wind power has an average NPV of 1.82 billion USD and 113.8 kilotonnes of emissions. The same solution with only gas power available has an average NPV of 0.07 billion USD and 197.9 kilotonnes of emissions. As seen in the lower row of Fig. 10, when only gas power is available, all controls approach zero, meaning no injection and production. The emissions do not approach zero here because of the constant baseload of 4MW that needs to be accounted for. The reason that not all of the controls are exactly zero could be related to inaccurate gradients due to the truncation of control samples that fall outside the bound of the domain. When wind power is available, the production controls approach the same solution. However, we still observe significantly more water injection compared to the solution with only gas. This could be a combination of two things: the gradient estimate is poor because many of the production controls are zero, and in the limit when $\omega = 1.0$ there is no mechanism in the optimization algorithm to prevent injection of water based on wind energy. However, the solutions we are the most interested in are on the other side of the spectrum for $\omega \leq 0.75$.


On the other side of the spectrum, when $\omega = 0.0$, the objective function only contains the NPV term. Here, the average NPV of the solution with wind power is 27.3 billion USD, and when only gas is available it is 26.7 billion USD. This is a difference of 600 million USD, most of which comes from extra gas that can be sold/exported rather than used as fuel in the gas turbines because more of the power demand is covered by wind power. The average amount of CO₂ emitted from the solution with wind is 223.7 kilotonnes, which is 116.2 kilotonnes less than the solution with only gas power. Using a cost of 0.15 USD/kg for emissions, this results in an average of only 17.4 million USD in saved CO₂ tax.

The Norwegian government, together with the oil and gas sector, has set a goal of a 40% reduction in emissions by the year 2030 compared to the emissions from 2005, and a net-zero goal for 2050 [24]. Considering the solution with no wind when $\omega = 0.0$ as a reference, where the average emissions are about 340 kilotonnes, a 40% reduction of this is about 204 kilotonnes. Looking at the result from Table 3 (or Fig. 9), the only solution without wind that can reach this goal is when $\omega = 1.0$, which is not a viable solution, as it corresponds to no production. However, with wind power, the solution when $\omega = 0.5$ has an average emission of around 200 kilotonnes, just below the reduction goal of 40%. Additionally, this solution has an average NPV of 26.04 billion USD, which is only a 2.3% reduction from the average reference NPV of 26.66. A reduction of 140 kilotonnes of emissions corresponds to only 21 million in saved CO₂ tax. This shows that a CO₂ tax alone is not enough to incentivize the goals of emission reduction. For comparison, the production rate of oil, gas, and water, as well as the emission rate of these two solutions, are shown in Fig. 11.

6 Summary

This paper presents a methodology for incorporating wind power into the optimization of control rates in reservoir management. Optimization is performed over an ensemble of realistic wind power series and geological realizations, to address the high uncertainty in wind power availability and the uncertainty of geological knowledge. The paper proposes a mathematical model for simulating daily wind speeds and uses Bayesian inference to fit the model to wind speed data from the North Sea. The resulting wind speed model is used to create a wind speed ensemble, which is then transformed into a power ensemble for optimization.

Fig. 11 A comparison of the oil, gas, water, and emissions rate of the two solutions $\omega = 0.0$ with only gas power (gray) and $\omega = 0.5$ with wind and gas power combined



The methodology is applied to a bi-objective optimization problem, where the two objective functions are emissions and net present value. The weighted sum method of multi-objective optimization is employed to find two sets of Pareto solutions. For the first set, the wind power is combined with gas power, while the second set is optimized with only gas power available for comparison. The key takeaways from the numerical experiments are as follows:

- When only the NPV term is present ($\omega = 0.0$), the average NPV using both gas and wind power reaches 27.3 billion USD. This is 0.6 billion USD higher than the case when only gas power is available. Simultaneously, the case with wind reduces the emissions by 116 kilotonnes from 340 kilotonnes to 224 kilotonnes. This results in 17.4 million USD being saved CO₂ tax.
- When only the emission term is present ($\omega = 1.0$), both solutions approach the strategy of zero production. However, because of the constant baseload, the emissions never reach zero.
- Using the emission from the solution with only gas power when only NPV is prioritized ($\omega = 0.0$) as a reference, it is not possible to reduce the emissions by 40% (which is the goal of the Norwegian oil and gas industry by the year 2030). The solution with wind power when $\omega = 0.5$, reaches this goal with only a 2.3% reduction in NPV.
- CO₂ tax alone is not enough to incentivize the oil and gas industry to reduce its carbon footprint.

In conclusion, optimizing over an ensemble of realistic time series for wind power and geological uncertainty can provide valuable information for field operators and decision-makers about future emissions and profit uncertainty. This is especially valuable when applied to bi-objective optimization, as in this paper, because it quantitatively measures the trade-off between the two. It is worth mentioning that the results shown in this paper are specific to the test-case setup presented here. Therefore, it is interesting to perform such an analysis on a more extensive and detailed setup for the reservoir in the future.

Acknowledgements The authors acknowledge funding from the Centre of Sustainable Subsurface Resources (CSSR), grant nr. 331841, supported by the Research Council of 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.

Funding Open access funding provided by University of Bergen (incl Haukeland University Hospital).

Data Availability Link to code: https://github.com/MathiasMNilsen/ECMOR2024_special_issue.git

Declarations

Conflicts of 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.

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. Norge, Offshore: Klima og miljørapport 2023. Technical report, Offshore Norge (2023)
2. Norwegian Offshore Directorat: Resource report 2019. Technical report, Norwegian Offshore Directorat (2019)
3. Equinor: Hywind Tampen. Equinor.no (2023)
4. Angga, I.G.A.G., Bellout, M., Kristoffersen, B.S., Bergmo, P.E.S., Slotte, P.A., Berg, C.F.: Effect of CO² tax on energy use in oil production: waterflooding optimization under different emission costs. *SN Appl. Sci.* **4**(11), 313 (2022). <https://doi.org/10.1007/s42452-022-05197-4>
5. Brouwer, D.R., Jansen, J.-D.: Dynamic optimization of waterflooding with smart wells using optimal control theory. *SPE J.* **9**(04), 391–402 (2004). <https://doi.org/10.2118/78278-PA>
6. 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>
7. Fonseca, 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>
8. Oliver, D., Sætrom, J., Skorstad, A., Raanes, P., Saksvik, T., Kolbjørnsen, O.: Robust multi-objective production optimization with co2 emissions reduction. In: European Conference on the Mathematics of Geological Reservoirs (2024). <https://doi.org/10.3997/2214-4609.202437098>
9. Oliver, D.S., Raanes, P.N., Skorstad, A., Sætrom, J.: Multi-objective reservoir production optimization: Minimizing co2 emissions and maximizing profitability. *Geoenergy Sci Eng.* **244**, 213396 (2025). <https://doi.org/10.1016/j.geoen.2024.213396>
10. Haakenstad, H., Breivik, Ø., Reistad, M., Aarnes, O.J.: NORA10EI: A revised regional atmosphere-wave hindcast for the North Sea, the Norwegian Sea and the Barents Sea. *Int. J. Clim.* **40**(10), 4347–4373 (2020). <https://doi.org/10.1002/joc.6458> <https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/joc.6458>
11. Brown, B.G., Katz, R.W., Murphy, A.H.: Time series models to simulate and forecast wind speed and wind power. *J. Appl. Meteorol. Clim.* **23**(8), 1184–1195 (1984) [https://doi.org/10.1175/1520-0450\(1984\)023<1184:TSMTSA>2.0.CO;2](https://doi.org/10.1175/1520-0450(1984)023<1184:TSMTSA>2.0.CO;2)
12. Oriol, A.-P., Virgile, A., Colin, C., Larry, D., Fonnesbeck, C.J., Maxim, K., Ravin, K., Jupeng, L., Luhmann, C.C., Martin, O.A., Michael, O., Ricardo, V., Thomas, W., Robert, Z.: PyMC: a modern and comprehensive probabilistic programming framework in Python. *PeerJ Comput. Sci.* **9**, 1516 (2023). <https://doi.org/10.7717/peerj-cs.1516>
13. Forde, N.: Forecasting with structural AR timeseries. In: PyMC Team (ed.) PyMC Examples. <https://doi.org/10.5281/zenodo.5654871>
14. Musial, W., Beiter, P., Smith, A., Tegen, S.: Potential Offshore Wind Energy Areas in California: An Assessment of Locations, Technology, and Costs. This study was funded by the U.S. Department of the Interior, Bureau of Ocean Energy Management. (2016). <https://doi.org/10.2172/1338174>
15. Jonkman, J., Butterfield, S., Musial, W., Scott, G.: Definition of a 5-mw reference wind turbine for offshore system development. Technical Report NREL/TP-500-38060, National Renewable Energy Laboratory (NREL) (2009). <https://docs.nrel.gov/docs/fy09osti/38060.pdf>
16. Liew, J., Urbán, A.M., Andersen, S.J.: Analytical model for the power-yaw sensitivity of wind turbines operating in full wake. *Wind Energy Sci.* **5**(1), 427–437 (2020). <https://doi.org/10.5194/wes-5-427-2020>
17. Goodarzi, E., Ziaei, M., Hosseinipour, E.Z.: Multiobjective Optimization, pp. 111–147. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-04400-2_4
18. Equinor: Conceptual description for Drogan (2020)
19. Zhou, X.-H., Wang, H., McClure, J., Chen, C., Xiao, H.: Inference of relative permeability curves in reservoir rocks with ensemble Kalman method. *Euro. Phys. J. E* **46**(6) (2023). <https://doi.org/10.1140/epje/s10189-023-00296-5>
20. Skjerve, K.B., Martinsen, F., Solbraa, E., Espelid, Ø.: ECalc - A Computationally Efficient Tool for Emission Forecasting. In: SPE Norway Subsurface Conference (2022). <https://doi.org/10.2118/209561-MS>

21. Rasmussen, A.F., Sandve, T.H., Bao, K., Lauser, A., Hove, J., Skaflestad, B., Klöfkorn, R., Blatt, M., Rustad, A.B., Sævareid, O., Lie, K.-A., Thune, A.: The Open Porous Media Flow reservoir simulator. *Comput. Math. Appl.* **81**, 159–185 (2021). <https://doi.org/10.1016/j.camwa.2020.05.014>
22. Fonseca, R.M., Kahrobaei, S.S., Gastel, L.J.T., Leeuwenburgh, O., Jansen, J.D.: Quantification of the impact of ensemble size on the quality of an ensemble gradient using principles of hypothesis testing. In: *SPE Reservoir Simulation Symposium*. SPE Reservoir Simulation Conference, pp. 031–011003 (2015). SPE-173236-MS
23. Nocedal, J., Wright, S.J.: *Numerical Optimization*, pp. 30–65. Springer, New York (2006). https://doi.org/10.1007/978-0-387-40065-5_3
24. Norge, Offshore: Framtidens energinæring på norsk sokkel. Technical report, Offshore Norge (2021)

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

Authors and Affiliations

Mathias M. Nilsen^{1,2} · Rolf J. Lorentzen¹ · Andreas S. Stordal¹ · Olwijn Leeuwenburgh³ · Eduardo Barros³

✉ Mathias M. Nilsen
mani@norceresearch.no

¹ NORCE Norwegian Research Centre, Bergen, Norway

² University of Bergen, Bergen, Norway

³ TNO, Utrecht, Netherlands