Preprints of the 22nd IFAC World Congress
Yokohama, Japan, July 9-14, 2023

Heteroscedastic Bayesian Optimisation for Active Power Control of Wind Farms ⋆

Kiet Tuan Hoang [*,***] **Sjoerd Boersma** [**] **Ali Mesbah** [***] **Lars Imsland** [*]

[*] *Department of Engineering Cybernetics, Norwegian University of Science and Technology, Norway. E-mail: kiet.t.hoang@ntnu.no, lars.imsland@ntnu.no*
[**] *Biometris Group, Wageningen University & Research, The Netherlands. E-mail: sjoerd.boersma@wur.nl*
[***] *Department of Chemical and Biomolecular Engineering, University of California, Berkeley, USA. E-mail: mesbah@berkeley.edu*

Abstract: Active power control of wind farms remains an open challenge due to inherent noise in wind power that arises from uncertain wind speed measurements and plant/model mismatch. To leverage the heteroscedastic nature of the wind power noise, heteroscedastic Bayesian optimisation (BO) is used for active power control of wind farms. BO utilises closed-loop performance data to tune the parameters of a stochastic model predictive controller (SMPC) in a systematic and data-efficient manner. This, in turn, allows for enhancing the closed-loop performance of the controller intended to decrease the power tracking error. A case study with 9 turbines in a 3x3 wind farm shows that the heteroscedastic BO approach achieves a reduced closed-loop power tracking error in terms of root-mean-square by 8.89% compared to one that relies on nominal BO and a decrease by 64.99% compared to a nominal model predictive controller (MPC) whose performance is not tuned using closed-loop data and BO.

Keywords: Wind farm; active power control; stochastic model predictive control; data-driven optimisation; Heteroscedastic noise; controller auto-tuning

1. INTRODUCTION

The development and deployment of wind turbines play a vital role in the transition to net-zero-emission societies in Europe (WindEurope, 2021). Wind turbines are commonly placed into grids to form wind farms such that maintenance and electricity cabling costs are reduced. The primary control objective of such farms is to reduce the total levelized cost of wind energy by efficient operation of the turbines. An important sub-problem for wind farm control is secondary frequency regulation (a sub-class of active power control), where the goal is to coordinate the operation of individual turbines to meet some desired power reference during a period of several minutes; for example, see Ela et al. (2014), Spudić et al. (2015), Jensen et al. (2016), and Siniscalchi-Minna et al. (2018).

However, most approaches for active power control do not explicitly consider the uncertainty inherent to wind farm control that can degrade control performance (Boersma et al., 2019a). This challenge can generally be addressed with an SMPC (Mesbah, 2016) to explicitly account for probabilistic uncertainties when their distributions can be quantified *a priori* (Boersma et al., 2019a). Nonetheless, quantifying uncertainties *a priori* can be impractical, as wind power noise is time-variant and *heteroscedastic* (i.e., input-dependent) due to uncertain measurement of the wind speed. Heteroscedasticity refers here to the property that the noise variance varies across some parameter (input) space. This can occur in power control of wind farms since turbines' power outputs are perturbed and dependent on time-varying, uncertain wind speed measurements through the controller. One approach to account and leverage for heteroscedastic noise in wind speed and, accordingly, wind power control is through probabilistic modeling with heteroscedastic Gaussian processes (GP) (Rogers et al., 2020). Yet, a better noise model does not necessarily translate to improved closed-loop performance.

This paper proposes a heteroscedastic BO (Guzman et al., 2021) approach to directly translate closed-loop data to closed-loop performance with controller auto-tuning (Khosravi et al., 2022, Paulson et al., 2022) on an SMPC controller for active power control. The heteroscedastic BO setup leverages closed-loop data directly to model the closed-loop cost for increased closed-loop performance by tuning the SMPC hyperparameters. BO provides a systematic and data-effective framework for utilising closed-loop data gathered from expensive high-fidelity simulations or actual experiments (Boersma et al., 2017) for SMPC tuning. Advantages of using BO involve the effective handling the high dimensional and nonlinear SMPC tuning parameters search space, which scales with the number of turbines in the farm.

⋆ The Research Council of Norway funds this research through PETROSENTER LowEmission (project code 296207), together with the American-Scandinavian Foundation and the Norway-America Association.

The paper is structured as follows. Section 2 presents the wind farm system under study. Section 3 presents the SMPC formulation for wind farms, along with the control-relevant model of the wind turbines. Section 4 presents the heteroscedastic BO approach for SMPC tuning using closed-loop data. The proposed method is validated in Section 5 for the case of uncertain wind speed measurements. Finally, concluding remarks are given in Section 6.

2. WIND FARM SYSTEM

This study focuses on a wind farm system from the SOWFA ALM simulations with 9 NREL 5 MW turbines in a 3x3 grid; see Fig. 1. Each of the turbines in the farm is controlled through axial induction-based control, where the power output P_i for turbine $i \in [1, N_t]$ is controlled by the disk-based thrust coefficients $C_i'(t)$ and yaw angle $\gamma_i(t)$ for N_t turbines. The wind speed $\mathbf{v}_i(t)$ is assumed measurable at each turbine i, where the measured wind speed $\hat{\mathrm{v}}_i(t) \sim \mathcal{N}\left(\mathrm{v}_i, \sigma^2_{\mathrm{v}_i}\right)$ is perturbed by normally-distributed noise. Active power control of such a farm aims at coordinating the turbines to minimise the tracking error $\mathbf{e}(t)$ by meeting some power reference $P^{\mathrm{ref}}(t)$

$$\mathbf{e}(t) = P^{\mathrm{ref}}(t) - \sum_{i=1}^{N_t} P_i(C_i'(t), \gamma_i(t), \hat{\mathrm{v}}(t), t), \tag{1}$$

while minimising excessive and oscillatory control inputs due to uncertainty in wind speed measurement.

3. STOCHASTIC MODEL PREDICTIVE CONTROL FOR WIND FARM CONTROL

This section first introduces a control-relevant model of a wind turbine for active wind power control. This is followed by formulating a sample-based SMPC to minimise the power tracking error (1) under uncertain wind speed measurements.

3.1 Control-relevant model of turbine

Boersma et al. (2019b) showed that axial induction-based wind farm power tracking could be achieved with the following model for turbine power output $P_i(t)$ and thrust coefficient $C_i'(t)$ while neglecting wake effects if the control objective is power tracking

$$P_i(t) = \frac{\pi D^2}{8} \left(\hat{\mathrm{v}}_i(t)\cos[\gamma_i(t)]\right)^3 \hat{C}_i'(t) \tag{2a}$$

$$C_i'(t) = \tau \frac{\hat{C}_i'(t)}{\mathrm{d}t} + \hat{C}_i'(t), \tag{2b}$$

where D is the rotor diameter and $\hat{C}_i'(t)$ is the first-order filtered wind turbine input that is applied to the plant. The wind farm model can then be modelled as N_t uncoupled subsystems consisting of (2a) and (2b). For control, (2a) and (2b) can be temporally discretised at sample period Δt using a zero-order hold method

$$\mathbf{x}_{i,k+1} = A_i \mathbf{x}_{i,k} + B_i(\hat{\mathrm{v}}_{i,k})\mathbf{u}_{i,k}, \tag{3}$$

where $\mathbf{x}_{i,k}^{\mathrm{T}} = [P_{i,k} \quad \hat{C}_{i,k}'] \in \mathbb{R}^2$, $\mathbf{u}_{i,k} = C_{i,k}' \in \mathbb{R}$, and

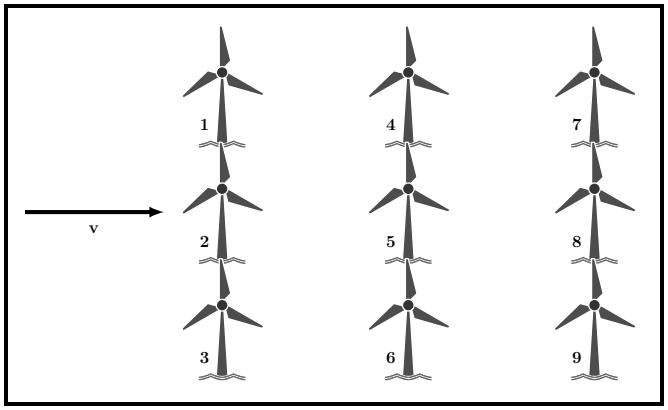

Fig. 1. An illustration of the chosen wind farm layout of a 3x3 grid with 9 turbines.

$$A_i = e^{-\Delta t/\tau} I_2 \in \mathbb{R}^{2x2} \tag{4a}$$

$$B_i = \int_0^{\Delta t} \frac{1}{\tau} e^{-s/\tau} \mathrm{d}s \begin{pmatrix} \begin{bmatrix} \frac{\pi D^2}{8} \hat{\mathrm{v}}_{i,k}^3 \\ 1 \end{bmatrix} \end{pmatrix} \in \mathbb{R}^2, \tag{4b}$$

with k being the current discrete time-step, and $\gamma_i = 0$.

3.2 Sample-based stochastic model predictive control

Commonly, an SMPC is formulated using an expectation-based operator $\mathbb{E}$ for the cost function $\mathbf{J}_{\mathrm{smpc}}$ with chance constraints defined in terms of the probability operator $\mathbb{P}$

$$\begin{aligned} \min_{\mathbf{x},\mathbf{u}} \quad & \mathbb{E}[\mathbf{J}_{\mathrm{smpc}}\left(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k\right)] \\ \text{s.t.} \quad & \mathbf{x}_{k+1} = f\left(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k\right) \\ & \mathbb{P}\left[h\left(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k\right) \leq 0\right] \leq 1 - \epsilon \\ & \mathbf{x_0} = \mathbf{x}_{k_0}, \quad \mathbf{w}_k \sim \mathcal{N}(\mu, \sigma^2) \\ & k = k_0, k_0+1, \ldots, k_0+N_P, \end{aligned} \tag{5}$$

where $\mathbf{x}_k$, $\mathbf{u}_k$, and $\mathbf{w}_k$ are the system states, inputs, and uncertainty, f is the state transition function, h are some inequality constraints, $\epsilon \in [0, 1]$ denotes the lower bound of the desired probability for which h should be satisfied under $\mathbf{w}_k$, $\mathbf{x}_{k_0}$ is the current measured state at time-step k_0, and N_P is the prediction horizon over which the SMPC optimises the control signals.

For wind farm power control, a tractable sample-based approximation of (5) without chance-constraints can be used with $\gamma_i = 0$, given that $\hat{\mathrm{v}}(t)$ can be sampled directly from $\mathcal{N}(\mathrm{v}, \sigma^2_{\mathrm{v}})$ while assuming that σ^2_{v} is known *a priori* (Boersma et al., 2019a)

$$\begin{aligned} \min_{\mathbf{C}_\mathbf{k}'} \quad & \frac{1}{N_s} \sum_{j=1}^{N_s} \sum_{k=k_0}^{k_0+N_p} \left(e_{j,k}^{\mathrm{T}} Q e_{j,k} + \sum_{i=0}^{N_t} \Delta C_{i,k}' R \Delta C_{i,k}' \right) \\ \text{s.t.} \quad & \mathbf{x}_{i,j,k+1} = A\mathbf{x}_{i,j,k} + B(\hat{\mathrm{v}}_{i,j,k})\mathbf{u}_{i,k} \\ & C_{\min}' \leq C_{i,k}' \leq C_{\max}' \\ & \mathbf{x_0} = \mathbf{x}_{k_0}, \quad \hat{\mathbf{v}}_{i,j,k} \sim \mathcal{N}(\mathrm{v}_{i,k}, \sigma^2_{\mathrm{v}_{i,k}}) \\ & k = k_0, k_0+1, \ldots, k_0+N_P \\ & j = 1, 2, \ldots, N_s, \quad i = 0, 1, \ldots, N_t, \end{aligned} \tag{6}$$

where N_s is the number of wind speed measurement samples $\hat{\mathrm{v}}_{i,j,k}$ from $\mathcal{N}(\mathrm{v}_{i,k}, \sigma^2_{\mathrm{v}_{i,k}})$, R and Q are constant

matrices to be tuned, $\Delta C'_{i,k} = (C'_{i,k} - C'_{i,k-1})$ is the input rate, and $\mathbf{e}_{j,k}$ is the power tracking error (1). The main idea in the sample-based SMPC (6) is that given a measured wind speed $\hat{\mathrm{v}}_i(t) \neq \mathrm{v}_i(t)$, N_s samples of $\hat{\mathrm{v}}_i(t)$ are chosen from $\mathcal{N}(\mathrm{v}_i, \sigma^2_{\mathrm{v}_i})$ to approximate $\hat{\mathrm{v}}_i(t) \approx \mathrm{v}_i(t)$. The sample-based SMPC increases tracking performance by reducing the oscillatory control around the reference. Similarly to nominal MPC, wind farm control is achieved by solving (6) at each sampling time k_0. Based on the receding-horizon principle, a control law κ_{smpc} for wind farm power control can be defined implicitly as

$$\kappa_{\mathrm{smpc}} = C'^{*}_{k0}(\mathbf{x}_{k_0}, \hat{\mathrm{v}}, P^{\mathrm{ref}}; \theta), \tag{7}$$

with $\theta \in \mathbb{R}^{N_t}$ comprising of σ_{v_i}.

However, such an SMPC can be hard to implement in practice, as sampling $\hat{\mathrm{v}}_i(t)$ directly from $\mathcal{N}(\mathrm{v}_i, \sigma^2_{\mathrm{v}_i})$ can be impractical since $\sigma^2_{\mathrm{v}_i}$ is time-varying and not known *a priori*. The sample-based SMPC must thus rely on $\hat{\mathrm{v}}_i(t)$ and sample from a distribution with N_t unknown covariances $\sigma^2_{\hat{\mathrm{v}}_i}$ to approximate $\mathrm{v}_i(t)$. Additionally, plant-model mismatch and the fact that the SMPC samples around $\hat{\mathrm{v}}_i(t)$ may result in an offset (Boersma et al., 2019a). For offset free tracking, an offset term $\alpha_{\mathbf{e}} \in [0, \infty]$ can be used to modify (1) accordingly

$$\mathbf{e} = P^{\mathrm{ref}} - \sum_{i=1}^{N_t} P_i + \alpha_{\mathbf{e}}. \tag{8}$$

Still, there is an open question on how to derive $\theta \in \mathbb{R}^{N_t+1}$ which comprises of $\sigma^2_{\hat{\mathrm{v}}_i}$ and $\alpha_{\mathbf{e}}$ such that the controller meets the reference, where the choice of $\sigma^2_{\hat{\mathrm{v}}_i}$ influences the output P_i through (7).

4. HETEROSCEDASTIC BAYESIAN OPTIMISATION FOR SMPC TUNING USING CLOSED-LOOP DATA

Deployment of SMPC for active power control of wind farms is challenging in practice, as it is non-trivial to specify the N_t unknown covariances $\sigma^2_{\hat{\mathrm{v}}_i}$ and $\alpha_{\mathbf{e}}$. The search space is high dimensional and nonlinear and scales with the number of turbines N_t. For example, a random grid approach with 9 turbines for the max, min, and the mean value results in 3^{N_t+1} combinations of parameters. Exploring such a parameter space is impractical, mainly due to the considerable computational cost of closed-loop high-fidelity simulations with the sample-based SMPC controller (Boersma et al., 2017). To this end, this section introduces a heteroscedastic BO approach for systematic and data-efficient selection of the parameters $\sigma^2_{\hat{\mathrm{v}}_i}$ and $\alpha_{\mathbf{e}}$ using closed-loop data that can quantify the closed-loop performance of the SMPC. For brevity, the identifier for each turbine i has been omitted in this section.

4.1 Quantifying closed-loop performance

With the implicit control law (7), a closed-loop system can be defined as

$$\mathbf{x}_{k+1} = f_{\mathrm{hf}}\left(\mathbf{x}_k, \kappa_{\mathrm{smpc}}(\mathbf{x}_k, \hat{\mathrm{v}}_k, P^{\mathrm{ref}}_k; \theta), \mathrm{v}_k\right), \tag{9}$$

where f_{hf} represents the dynamics of the wind farm, available through a high-fidelity simulator or experiment that can be queried. A closed-loop trajectory of the wind farm can be defined in terms of θ and some set of uncertain variables $\mathcal{W} = \{\mathbf{x}_0, \mathrm{v}_0, \mathrm{v}_1, \ldots, \mathrm{v}_{N-1}\}$ by recursively applying (9) N times to compute a closed-loop trajectory with $\mathbf{x}_k = \kappa_{\mathrm{smpc}}(\mathbf{x}_{k-1}, \hat{\mathrm{v}}_{k-1}, P^{\mathrm{ref}}_{k-1}; \theta)$

$$z_{cl}(\theta, \mathcal{W}) = \{\theta, \mathbf{x}_0, \mathrm{v}_0, \ldots, \mathbf{x}_{N-1}, \mathrm{v}_{N-1}, \mathbf{x}_N\}. \tag{10}$$

BO relies on quantifying the closed-loop performance of the SMPC, based on the controller parameters that are tuned. A natural candidate for active power control is the root-mean-square of the closed-loop power tracking error with $\mathbf{e}_{\mathrm{cl}}$ as defined according to (8),

$$\mathbf{J}_{\mathrm{cl}}(z_{cl}(\theta, \mathcal{W})) = \mathbb{E}\left[\sqrt{\frac{1}{N}|\mathbf{e}_{\mathrm{cl}}|^2}\right] = \mathbb{E}\left[\mathbf{e}_{\mathrm{cl,rms}}\right]. \tag{11}$$

This closed-loop performance measurement is inherently noisy due to the uncertain wind speed measurements $\hat{\mathrm{v}}_k$ and the mismatch between the high-fidelity simulator f_{hf} and the control-relevant model (2). The plant-model mismatch is difficult to quantify in terms of noise, but $\hat{\mathrm{v}}_k$ is propagated through (2a) with (7)

$$\hat{P}_{\mathrm{cl}}(\theta) = \frac{\pi D^2}{8}\hat{\mathrm{v}}^3 \kappa_{\mathrm{smpc}}(\hat{\mathbf{x}}, \hat{\mathrm{v}}, P^{\mathrm{ref}}; \theta), \tag{12}$$

which exhibits heteroscedastic properties as the variance from (12) depends on the parameter space of θ. The residue can be computed as the root-mean-square of (12) and the actual closed-loop wind farm power output P_{cl}

$$\epsilon_{\mathrm{cl}}(z_{cl}(\theta, \mathcal{W})) = \sqrt{\frac{1}{N}\left|\hat{P}_{\mathrm{cl}} - P_{\mathrm{cl}}\right|^2}. \tag{13}$$

4.2 SMPC tuning using closed-loop performance data

The SMPC parameters $\theta = \{\sigma_{\hat{v}}, \alpha_{\mathbf{e}}\}$ can now be derived by optimising (11)

$$\theta^* = \arg\min_{\theta} \ \mathbf{J}_{\mathrm{cl}}(z_{cl}(\theta, \mathcal{W})), \tag{14}$$

with BO that uses a probabilistic surrogate model of the closed-loop performance measure $\mathbf{J}_{\mathrm{cl}}$ for exploring the search space of θ (Shahriari et al., 2016). Generally, BO can be applied to objectives that do not have a closed-loop functional form, such as (14), by modeling the closed-loop performance using GP regression

$$y = \mathbf{J}_{\mathrm{cl}}(z_{cl}(\theta, \mathcal{W})) + \epsilon_{\mathrm{cl}} \sim \mathrm{GP}(m(\theta), k(\theta, \theta')), \tag{15}$$

where $\epsilon_{\mathrm{cl}} \sim \mathcal{N}(0, \sigma^2_\epsilon)$ is a zero-mean normal random variable, and the GP is specified by its mean function $m(\theta)$ and a positive-definite covariance function $k(\theta, \theta')$ (Rasmussen, 2005). The main idea with BO then is to utilise the GP from (15) to compute the posterior distribution of $\mathbf{J}_{\mathrm{cl}}(z_{cl}(\theta, \mathcal{W}))$, which is used in an acquisition function to quantify the trade-off between exploring new sets of θ in the search space which BO has not considered yet and exploiting current sets of θ. This exploration and exploitation are based on the uncertainty description provided by GP. The regular BO framework can be extended to include heteroscedasticity when σ_ϵ is not constant and varies as a function of θ. This is achieved by letting σ_ϵ be defined by

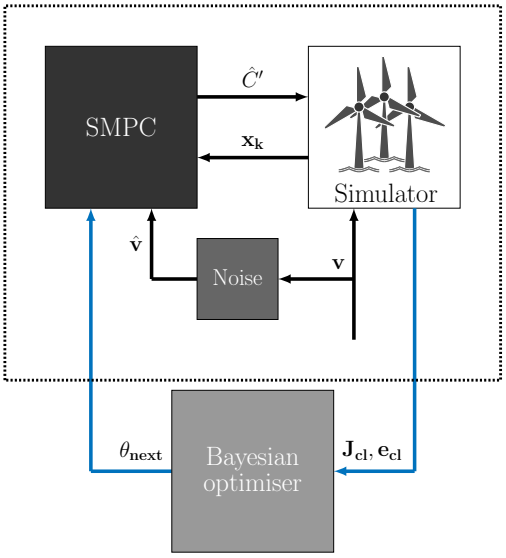

Fig. 2. A schematic of the heteroscedastic BO approach for auto-tuning the SMPC.

another GP to model the log of σ_ϵ (Goldberg et al., 1998)

$$\log(\sigma_\epsilon(\theta)) \sim \mathrm{GP}_\epsilon\left(0, \left(k_\epsilon\left(\theta, \theta'\right)\right)\right). \tag{16}$$

A schematic of the heteroscedastic BO approach for auto-tuning of SMPC is given in Fig. 2. Here, the noise block is an unknown function for which the actual wind speed is perturbed by noisy wind speed measurement $\hat{\mathbf{v}}(\mathbf{t}) \neq \mathbf{v}(t)$. For every closed-loop simulation, BO tunes the parameters θ of SMPC using closed-loop performance data. BO continues to gather closed-loop data until the optimal values θ^* are found or the computational budget has been reached.

5. SIMULATION STUDY

In this section, a nominal model predictive controller (MPC) is compared to a stochastic model predictive controller (SMPC), which is tuned using both homoscedastic and heteroscedastic BO from closed-loop simulations. In the simulations, (6) is used in both the MPC and the SMPC for optimal control, except that the MPC is initialised with $(N_s, \alpha_{\mathbf{e}}, \sigma_{\hat{\mathrm{v}}_i}) = (1, 0, 0),\ \forall i \in [1, N_t]$. Additionally, although unrealistic, an MPC with perfect knowledge of the wind speed $\mathrm{v}(t)$ is also included ($\hat{\mathrm{v}}(t) = \mathrm{v}(t)$). For BO, the heteroscedastic BO includes (16) to model time-varying σ_ϵ in the GP, which the nominal (homoscedastic) BO omits as it assumes time-invariant σ_ϵ.

5.1 Simulation environment and variables

A computer with an Intel(R) Core(TM) i7-9850H CPU @ 2.60 GHz is used for this simulation study. The controller and the simulations are handled in MATLAB with YALMIP (Löfberg, 2004) and MOSEK (MOSEK, 2019) as the solver for the predictive controller and WFSim (Boersma et al., 2018) as the simulator of choice to test the proposed method on. The resulting predictive controller is tuned from closed-loop data from WFSim with BO using Ax (Bakshy et al., 2018) and BoTorch (Balandat et al., 2020) in Python.

Table 1. Wind farm layout constants

Symbol	Value
$L_x \times L_y$	1.9×0.8 [km]
$\Delta x \times \Delta y$	28×38 [m]
D, z_t	126.4, 90 [m]
$\Delta t, N$	1, 998 [s]
N_t	9 [-]

Every simulation is initialised with a fully developed wake. This means that the flow has been simulated beforehand with $(C_i', \gamma_i) = (2, 0)$, where the value for C_i' corresponds to the Betz-optimal value. Initial free-stream wind speeds are defined according to $(U_\infty, V_\infty) = (12, 0)$ [m/s], where U_∞, V_∞ are the boundary wind speeds in the x- and y-direction, respectively. Further, site-dependent constants relevant to the simulation setup can be found in Table 1, where Δt is chosen such that the Courant condition holds (Courant et al., 1967), N is the simulation time, D is the rotor diameter, z_t is the hub-height of each turbine, Δ_x, Δ_y denotes the grid dimensions in x- and y-direction, and L_x, L_y denotes the length in x-, y-direction.

The power reference to be met is defined in this study through the greedy power P^{greedy}, which is the time-averaged power given $(C_i', \gamma_i) = (2, 0)$. Following P^{greedy}, a reference P_k^{ref} at each discrete step k can be defined

$$P_{\mathbf{k}}^{\mathrm{ref}} = 0.9P^{\mathrm{greedy}} + 0.2P^{\mathrm{greedy}}\delta P_k, \tag{17}$$

where δP_k follows a normalised automatic generation control (AGC) signal (Pilong, 2013) and $P^{\mathrm{greedy}} = 34.4$ [MW]. Given (17), the performance measure for the subsequent simulation study will follow the root-mean-square of the power tracking error; see (11).

This study assumes that the generated power output $P_i(t)$ and perturbed average rotor wind speed $\hat{\mathrm{v}}_i$ for each of the turbine i in WFSim are measured directly, where $\hat{\mathrm{v}}_i(t) \sim \mathcal{N}\left(\mathrm{v}_i + \mathbf{w}_{\mathrm{sensor}}, \sigma_{\mathrm{v}_i}^2\right)$, with $\mathbf{w}_{i,\mathrm{sensor}} \sim \mathcal{N}\left(0, \sigma_{i,\mathrm{sensor}}^2\right)$ being the uncertainty associated with each turbine and σ_{v_i} follows the magnitude of the wind speed

$$\sigma_{\mathrm{v}_i} = \frac{\mathrm{v}_i^2}{K_\sigma}, \tag{18}$$

where K_σ is some constant. The idea behind an individual sensor noise is that no sensors are the same, while (18) is chosen in this simulation study to model the wind speed uncertainty as noise with a constant covariance is non-realistic.

The BO with both heteroscedastic and homoscedastic GPs is set to run with a limited budget of 100 trials, as each simulation is expensive. Each trial and BO method are implemented with a Monte-Carlo-based batch noisy expected improvement acquisition function. One parameter $\sigma_{\hat{\mathrm{v}}_i}$ is derived for each turbine i, which together with $\alpha_{\mathbf{e}}$ results in

Table 2. Controller constants

Symbol	Value
Q, R	$10^{\ 4}$,10^9 [-]
$\Delta t, N_p, \tau$	1, 10, 5 [s]
$\mathbb{U}$	[0.1, 2] [-]
N_s	20 [-]

10 parameters, which are initialised to be zero. The SMPC and MPC parameters are given in Table 2 unless stated otherwise, where the initial system state is computed by simulating WFSim once in open-loop given the configurations in `layoutSet_sowfa_9turb_apc_alm_turbl`.

For clarification, $\mathbf{J}_{\text{cl}}$ (11) is used interchangeably with the power tracking error $\mathbf{e}_{\text{rms}}$ in the subsequent analysis.

5.2 9-turbine case study

It can be interesting to start by investigating the noise profile expected from wind farm control. As such, 150 closed-loop simulations are performed with the SMPC for $\sigma_{\hat{\text{v}}_i} = 2,\ \forall i \in [1, N_t]$, while varying α_{e}. Fig. 3 shows the resulting residue plot of $\log(\mathbf{J}_{\text{cl}})$ where a blue line has been included to clearly show the origin, while the red dots are the residues. The figure shows a clear trend, where the spread of the residue of $\log(\mathbf{J}_{\text{cl}})$ increases with $\alpha_{\mathbf{e}}$. Such a trend can also be established for $\sigma_{\hat{\text{v}}_i},\ \forall i \in [1, N_t]$ but is not shown here for brevity.

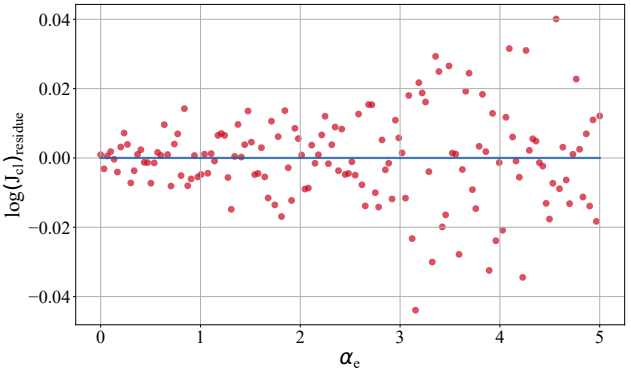

Fig. 3. A plot showing the residue of the log of the closed-loop performance measure $\mathbf{J}_{\text{cl}}$ against α_{e} with $\sigma_{\hat{\text{v}}_i} = 2,\ \forall i \in [1, N_t]$.

Now that heteroscedasticity has been investigated, a comparison between the convergence and the performance between a BO with a prior over the residue (heteroscedastic BO) compared to a nominal BO (homoscedastic BO) in terms of convergence, and performance is warranted. As such, Fig. 4 shows the mean tracking error as $\mathbf{J}_{\text{cl}}$ plotted against each BO iteration after restarting and applying each method 5 times, which results in 5 different optimal parameter configurations for each of the methods. As shown in Fig. 4, the performance of the heteroscedastic BO is, on average, higher than the homoscedastic BO, which does not converge.

The results are verified with 50 closed-loop Monte-Carlo simulations for each of the 5 different optimal parameter configurations for each method for different realisations of $\hat{\text{v}}(t)$ in Fig. 5, which shows the mean time profiles of the combined wind turbines for an SMPC with derived parameters θ from heteroscedastic BO as blue lines, an SMPC with homoscedastic BO as red lines, an MPC where measured wind speed $\hat{\text{v}}(t)$ is the same as the actual wind speed $\text{v}(t)$ as green lines, and an MPC where $\hat{\text{v}}(t) \neq \text{v}(t)$ as yellow lines. The MPC where $\hat{\text{v}}(t) = \text{v}(t)$ acts here as the unrealistic benchmark with the MPC where $\hat{\text{v}}(t) \neq \text{v}(t)$ as the realistic case.

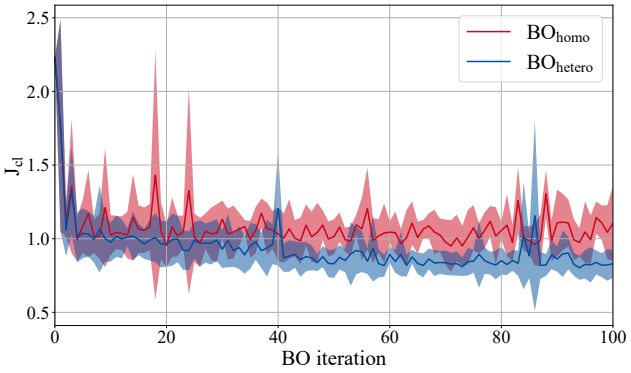

Fig. 4. A plot showing the mean power tracking error as $\mathbf{J}_{\text{cl}}$ against the current Bayesian optimisation trial. The shaded areas are the average trajectories with ± 1 standard deviation.

Fig. 5 shows that the SMPC with heteroscedastic BO and homoscedastic BO achieves a comparable performance in comparison with the benchmark MPC where $\hat{\text{v}}(t) = \text{v}(t)$, where the former achieves better tracking between the two. Some oscillations around the reference are expected. In contrast, the realistic MPC where $\hat{\text{v}}(t) \neq \text{v}(t)$ suffers greatly in tracking performance, both due to plant-model mismatch but also due to uncertainty in wind speed measurement. These results can be confirmed by looking at Table 3 which shows the mean tracking error $\mathbf{e}_{\text{rms}}$ for each method with their standard deviation.

Table 3. Power tracking error mean and standard deviations

Method	$\mu(\mathbf{e}_{\text{rms}})$ [MWH]	$\sigma(\mathbf{e}_{\text{rms}})$ [MWH]
SMPC + $\mathrm{BO_{homo}}$	0.8602	0.0144
SMPC + $\mathrm{BO_{hetero}}$	0.7837	0.0129
MPC ($\hat{\text{v}} = \text{v}$)	0.5514	0.0000
MPC ($\hat{\text{v}} \neq \text{v}$)	2.2389	0.0145

As expected, the mean μ and standard deviation σ of $\mathbf{e}_{\text{rms}}$ is the lowest for when v is used directly in an MPC, which in practice is unrealistic. The second best in terms of $\mathbf{e}_{\text{rms}}$ is with the SMPC + $\mathrm{BO_{hetero}}$, with the SMPC + $\mathrm{BO_{homo}}$ lagging in terms of performance. Lastly, the

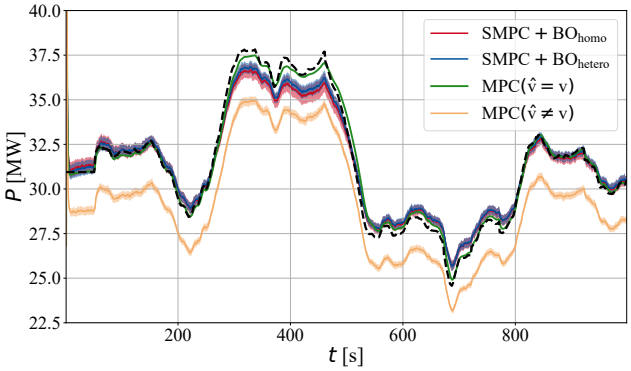

Fig. 5. Time profiles of the total wind farm power output mean for the different methods. The reference to be followed is illustrated as black stipulated lines, and the shaded areas are the average trajectories with ± 1 standard deviation.

nominal MPC using measured wind $\hat{v}$ directly has the highest $\mathbf{e}_{\text{rms}}$. Numerically this means that an SMPC + $\text{BO}_{\text{hetero}}$ approach yields on average a 64.99 % decrease in e_{rms} compared to the realistic case with an MPC ($\hat{v} \neq v$), and an 8.89% decrease compared to an SMPC + BO_{homo} approach. The increase in performance between SMPC + $\text{BO}_{\text{hetero}}$ over SMPC + BO_{homo} is expected as the former converges for the specified budget of 100 trials, whereas the latter does not. This performance difference can also be seen from Fig. 5 where the tracking performance of SMPC + $\text{BO}_{\text{hetero}}$ is on average better than the one from SMPC + BO_{homo}. The effect of applying heteroscedastic BO over a nominal BO can especially be seen in Fig. 5 between 0 [s] and 100 [s] where the plant-model mismatch is negligible as the reference and ideal MPC ($\hat{\mathbf{v}} = \mathbf{v}$) are almost the same. It can therefore be inferred that this area is mostly affected by the heteroscedastic noise from uncertain wind measurement, where using a heteroscedastic BO yield increased performance.

6. CONCLUDING REMARKS

This study demonstrates how heteroscedastic BO can be leveraged for increased wind farm power tracking performance, given heteroscedasticity in wind power noise. Results show an average decrease in tracking error by 8.89 % compared to one which relies on nominal BO when used together with an SMPC and a reduction by 64.99 % compared to just a nominal MPC without BO for a 9-turbine wind farm case study. Future work should investigate other ways uncertainty can impact wind farm control. For example, an economic formulation is generally required in wind farms to meet the power demand while at the same time minimise, for example, degradation factors. Such applications call for multi-objective-based optimisation, which BO also readily offers.

REFERENCES

E. Bakshy, L. Dworkin, B. Karrer, K. Kashin, B. Letham, A. Murthy, and S. Singh. Ae: A domain-agnostic platform for adaptive experimentation. 2018.

M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In *Advances in Neural Information Processing Systems 33*, 2020.

S. Boersma, B. Doekemeijer, P. Gebraad, P. Fleming, J. Annoni, A. Scholbrock, J. Frederik, and J.-W. van Wingerden. A tutorial on control-oriented modeling and control of wind farms. In *American Control Conference (ACC)*, pages 1–18, 2017.

S. Boersma, B. Doekemeijer, M. Vali, J. Meyers, and J.-W. van Wingerden. A control-oriented dynamic wind farm model: Wfsim. *Wind Energy Science*, 3(1):75–95, 2018.

S. Boersma, B. Doekemeijer, T. Keviczky, and J. van Wingerdenl. Stochastic model predictive control: uncertainty impact on wind farm power tracking. In *American Control Conference (ACC)*, pages 4167–4172, 2019a.

S. Boersma, B. Doekemeijer, S. Siniscalchi-Minna, and J. van Wingerden. A constrained wind farm controller providing secondary frequency regulation: An LES study. *Renewable Energy*, 134(C):639–652, 2019b.

R. Courant, K. Friedrichs, and H. Lewy. On the partial difference equations of mathematical physics. *IBM Journal of Research and Development*, 11(2):215–234, 1967.

E. Ela, V. Gevorgian, P. Fleming, Y. C. Zhang, M. Singh, E. Muljadi, A. Scholbrook, J. Aho, A. Buckspan, L. Pao, V. Singhvi, A. Tuohy, P. Pourbeik, D. Brooks, and N. Bhatt. Active power controls from wind power: Bridging the gaps. 2014.

P. Goldberg, C. Williams, and C. Bishop. Regression with input-dependent noise: A Gaussian process treatment. *Advances in Neural Information Processing Systems*, 10, 02 1998.

R. Guzman, R. Oliveira, and F. Ramos. Heteroscedastic Bayesian optimisation for stochastic model predictive control. *IEEE Robotics and Automation Letters*, 6(1): 56–63, 2021.

T. N. Jensen, T. Knudsen, and T. Bak. Fatigue minimising power reference control of a de-rated wind farm. *Journal of Physics: Conference Series*, 753(5):052022, 2016.

M. Khosravi, V. N. Behrunani, P. Myszkorowski, R. S. Smith, A. Rupenyan, and J. Lygeros. Performance-driven cascade controller tuning with Bayesian optimization. *IEEE Transactions on Industrial Electronics*, 69(1):1032–1042, 2022.

J. Löfberg. Yalmip : A toolbox for modeling and optimization in matlab. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.

A. Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016.

A. MOSEK. *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019.

J. Paulson, G. Makrigiorgos, and A. Mesbah. Adversarially robust Bayesian optimization for efficient auto-tuning of generic control structures under uncertainty. *AIChE Journal*, 68, 01 2022.

C. Pilong. PJM Manual 12: Balancing Operations. *technical report, PJM*, 2013.

C. K. I. Rasmussen, C. E. & Williams. *Gaussian processes for machine learning. MIT Press.* 2005.

T. Rogers, P. Gardner, N. Dervilis, K. Worden, A. Maguire, E. Papatheou, and E. Cross. Probabilistic modelling of wind turbine power curves with application of heteroscedastic Gaussian process regression. *Renewable Energy*, 148:1124–1136, 2020.

B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

S. Siniscalchi-Minna, F. D. Bianchi, and C. Ocampo-Martinez. Predictive control of wind farms based on lexicographic minimizers for power reserve maximization. In *American Control Conference (ACC)*, pages 701–706, 2018.

V. Spudić, C. Conte, M. Baotić, and M. Morari. Cooperative distributed model predictive control for wind farms. *Optimal Control Applications and Methods*, 36(3):333–352, 2015.

WindEurope. *Wind energy in Europe: 2021 Statistics and the outlook for 2022-2026*, 2021.