13th International Symposium on Turbulence and Shear Flow Phenomena (TSFP13)
Montreal, Canada, June 25–28, 2024

OPTIMISING ACTIVE FLOW CONTROL STRATEGIES FOR RANDOM AND CONTROLLED WIND SPEEDS VIA BAYESIAN OPTIMISATION

Mike Diessner
School of Computing
Newcastle University
Newcastle upon Tyne NE4 5TG, UK
m.diessner2@newcastle.ac.uk

Xiaonan Chen
School of Engineering
Newcastle University
Newcastle upon Tyne NE1 7RU, UK
xiaonan.chen@newcastle.ac.uk

Kevin J. Wilson
School of Mathematics, Statistics and Physics
Newcastle University
Newcastle upon Tyne NE1 7RU, UK
kevin.wilson@newcastle.ac.uk

Richard D. Whalley
School of Engineering
Newcastle University
Newcastle upon Tyne NE1 7RU, UK
richard.whalley@newcastle.ac.uk

ABSTRACT

Most simulations and experiments that control the turbulent boundary layer to reduce the skin-friction drag, assume that all influential variables, such as the free-stream velocity and the blowing amplitude, can be controlled. In a real-world application, it is likely that some variables are given externally by the environment—such as the wind speed—and blowing strategies must be selected accordingly. This study brings the optimisation of blowing actuators closer to real-life conditions by enabling optimisation of randomly varying free-stream velocity. Bayesian optimisation is extended to dynamic environments with controllable and uncontrollable variables by fitting a global surrogate model over all variables but optimising only the controllable variables conditional on the environmental variables. By conditioning on measurements of the uncontrollable variables, optima for their full domain can be predicted. This is in contrast to keeping environmental variables fixed to a single value, where only one optimum is found and the experiment must be repeated multiple times for different values to achieve similar results. The presented approach increases the available information within a single optimisation run and results in a more sample-efficient and cost-effective algorithm. As an example application, the method is applied to a 5-dimensional wind farm simulator to maximise the energy production conditional on the wind speed by controlling the derating of five wind turbines. The new method outperforms the Nelder-Mead algorithm by 2.2–60.0% and performs comparably to standard Bayesian optimisation for five selected wind speeds while allowing predictions of optimal derating levels for the full range of wind speeds.

INTRODUCTION

Turbulent forces near a vehicle's surface, such as an aircraft, are responsible for over half of the vehicle's energy consumption. Reducing the turbulent skin-friction drag realises monetary benefits due to reduced fuel consumption, and environmental and public health benefits due to lower CO2 emissions. Even a small drag reduction of 3% would save £1 million in jet fuel costs annually per aircraft and, with around 26,000 aircrafts around the world, would have a large global impact (Bushnell & Hefner, 1990). While low-amplitude blowing appears to be a viable solution to reduce drag (Hwang, 2004; Kornilov & Boiko, 2012), it is still unclear what blowing strategies are optimal as experiments are complex and expensive to conduct, prohibiting an exhaustive search of all strategies and the use of most optimisation algorithms that rely on gradient information and a large number of evaluations.

Bayesian optimisation (Jones *et al.*, 1998; Snoek *et al.*, 2012; Frazier, 2018) has emerged as an exception as it is a sample-efficient optimisation strategy designed for expensive experiments and simulations. In recent years it has been used in the field of computational fluid dynamics to minimise skin-friction drag via active control of blowing actuators (Mahfoze *et al.*, 2019; Diessner *et al.*, 2022; O'Connor *et al.*, 2023; Mallor *et al.*, 2023). These studies assume a fully controllable environment where all variables can be set to any desired value. However, a more realistic assumption is that some variables can be controlled while others cannot. For example, consider an aircraft for which the drag can be reduced by controlling uniform low-amplitude blowing actuators. The optimal blowing amplitude will depend on the cruising speed of the aircraft. The speed cannot be controlled to reduce the skin-friction drag as it is set by external factors such as the aircraft's schedule, course and weather. The maximisation of drag reduction can only be viewed conditionally on the aircraft's cruising speed, a problem which has not been tackled by previous studies where the speed is assumed to be fixed. This assumption means that the experiment would require repetition for all possible travelling speeds—an impracticable if not infeasible task.

This study uses an extension of Bayesian optimisation to optimise problems within dynamic environments with changing conditions that cannot be controlled. The approach fits a global surrogate model over all variables but optimises only the controllable variables conditional on measurements of the uncontrollable variables (Krause & Ong, 2011; Diessner *et al.*, 2024). In this article, Bayesian optimisation and the extension to account for environmental conditions are introduced before an example application to a 5-dimensional wind farm simulator with the objective of maximising the mean energy production conditional on randomly changing wind speeds is discussed. Lastly, findings are summarised and an outlook of future work is given.

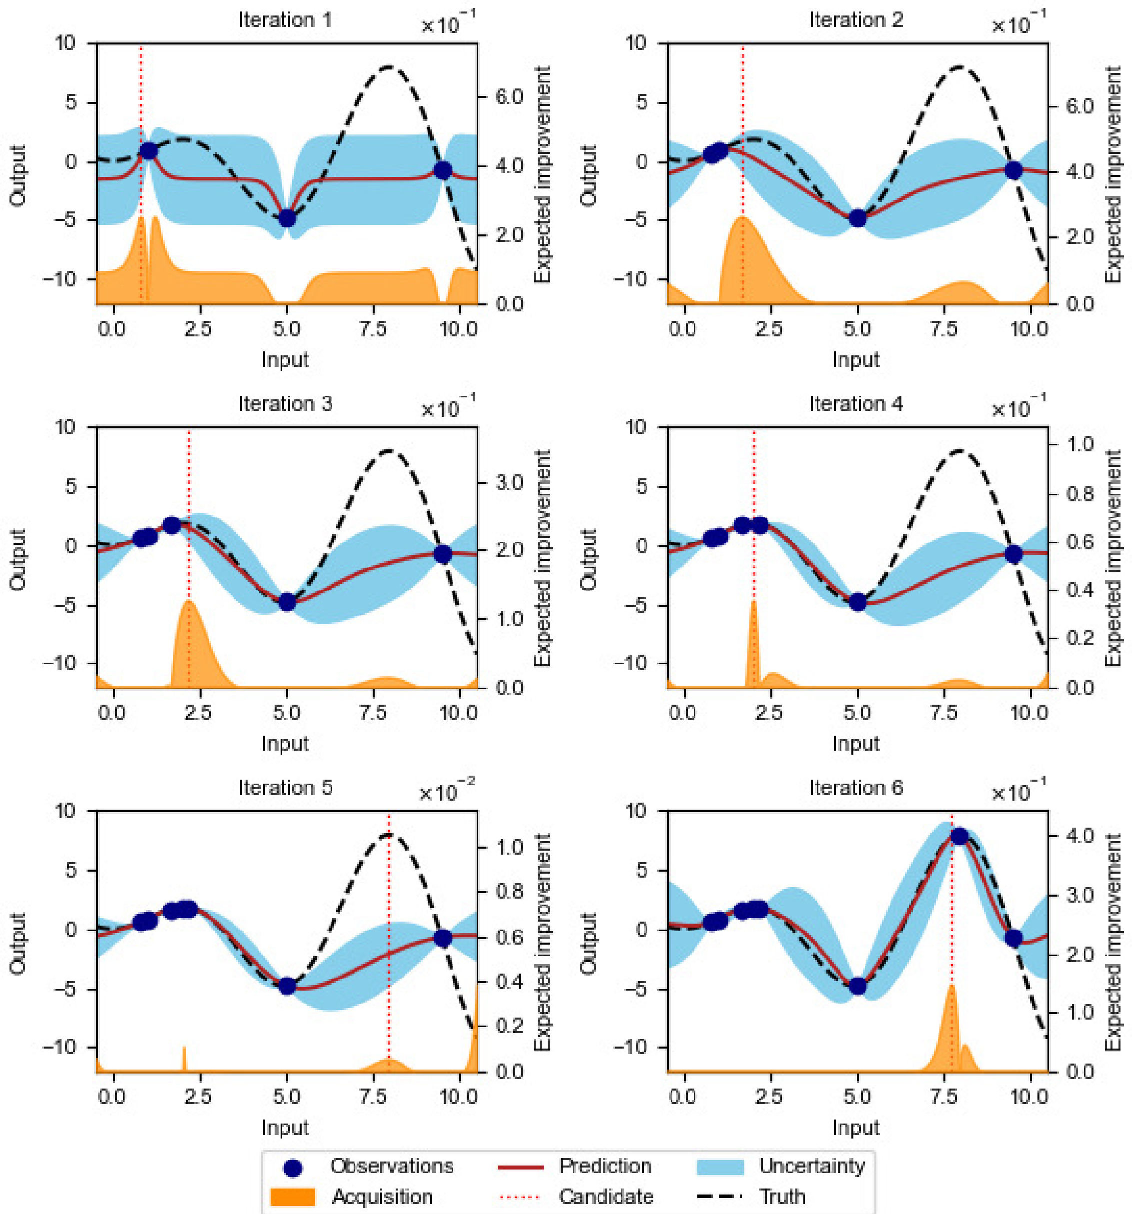

Figure 1: Bayesian optimisation applied to a 1-dimensional function with one local and one global maximum. Expected improvement is used as the acquisition function. The input space is bounded by $[0, 10]$.

BAYESIAN OPTIMISATION

Bayesian optimisation (Jones *et al.*, 1998; Snoek *et al.*, 2012; Frazier, 2018) is a sample-efficient and cost-effective optimisation strategy for expensive black-box functions such as physical experiments and computer simulations. It is performed in a loop where it sequentially selects a candidate point—a set of specific input variable values—at each iteration whose output is to be observed by conducting an experiment. Data points from previous iterations inform subsequent iterations ensuring that the algorithm learns continuously and explores the variable space effectively. The algorithm is stopped when an adequate solution is found or a specified evaluation budget is exhausted.

As the underlying objective function is unknown or too complex to solve analytically, a surrogate model is fitted to all available data to represent the best estimate of the objective function. Typically, a Gaussian process—a very flexible non-parametric model that can represent various objective functions (Rasmussen & Williams, 2006)—is chosen as the surrogate model. A Gaussian process provides a prediction and the corresponding uncertainty quantification for that prediction. Bayesian optimisation uses both to guide the sequential selection of candidate points by maximising an acquisition criterion. A popular acquisition criterion is expected improvement (Jones *et al.*, 1998) that selects candidate points with the highest probability of improving upon the data point with the best output value up to the present time.

The optimisation strategy is illustrated in Figure 1 on a 1-dimensional function with one local and one global maximum. In iteration 1, a Gaussian process is fitted to three observations that form the training data. The resulting prediction and uncertainty are used to compute the expected improvement acquisition criterion. Maximising this criterion yields the next candidate point that is to be observed from the true objective function. The candidate point is added to the other available observations which are used as the training data at it-

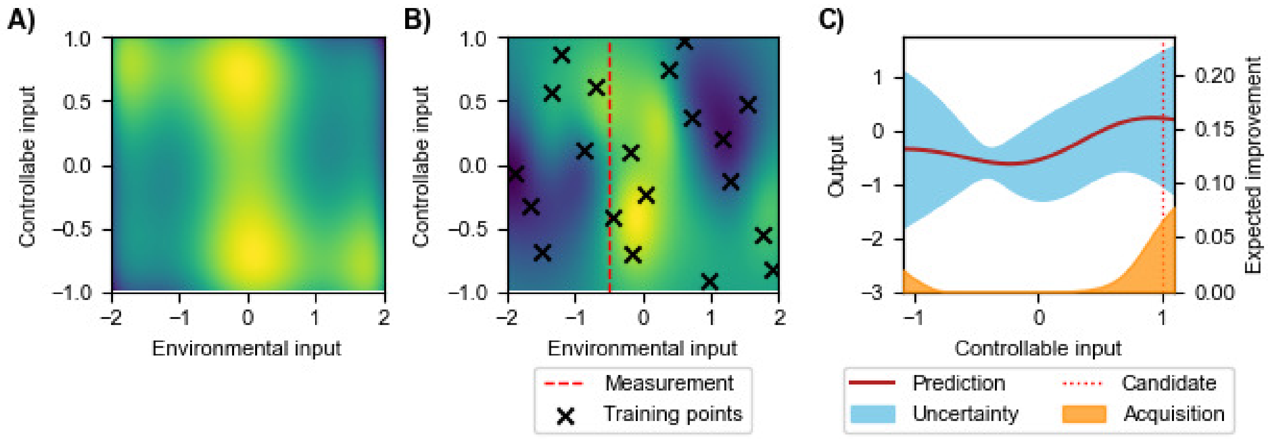

Figure 2: Optimisation of a 2-dimensional problem with one uncontrollable environmental variable x_1 and one controllable variable x_2. Yellow areas indicate highs and dark blue areas indicate lows. **A)** True objective function. **B)** Prediction of the surrogate model—a Gaussian process—with a measurement of the environmental variable taken for the next conditional optimisation step. **C)** Bayesian optimisation step conditional on the measurement of the environmental variable.

eration 2. With each iteration, Bayesian optimisation explores the variable space and the prediction of the Gaussian process represents the true objective function more accurately until the global maximum is found at iteration 5. The algorithm does not get stuck in the local maximum at iterations 3 and 4 due to the acquisition criterion that balances exploration—searching in areas with high Gaussian process uncertainty—and exploitation—searching in areas with high Gaussian process prediction. This study uses NUBO (Diessner *et al.*, 2023), a transparent open-source Bayesian optimisation framework written in Python, to optimise the simulations.

DYNAMIC ENVIRONMENTS

Consider a dynamic environment in which controllable and uncontrollable environmental variables affect the outcome of an experiment. As environmental variables cannot be controlled, the main objective of optimisation is to find optimal values for the controllable variables for any possible combination of environmental variable values. The standard Bayesian optimisation algorithm presented in the previous section requires three changes to allow optimisation conditional on the environmental variables (Diessner *et al.*, 2024). The resulting algorithm is referred to as ENVBO.

Firstly, standard Bayesian optimisation uses only controllable variables in its optimisation process and implicitly assumes that all other variables are fixed or not influential to the output. This assumption does not hold when optimising problems with dynamic environments where changing environmental variables affect the output of the objective function. Diessner *et al.* (2024) extend the surrogate model to include controllable and environmental variables. This guarantees that ENVBO uses all available information about the output of the objective function.

Secondly, ENVBO bases the acquisition function on the global surrogate model and requires a change to the maximisation strategy of the standard Bayesian optimisation approach as the environmental variables cannot be set to any arbitrary values that result from maximisation. Indeed, the maximisation of the acquisition function must depend on the environmental variable values given at the moment of maximisation. Thus, measurements for all environmental variables are taken after the surrogate model is fitted to the training data and the acquisition criterion is maximised conditionally on these measurements. Conditional optimisation means that the uncontrollable environmental variables are fixed at the taken measurements while optimal values for the controllable variables are found by maximising the acquisition criterion. The new candidate point is a combination of computed values for the controllable variables and measurements for the environmental variables. It is then evaluated by repeating the experiment.

Lastly, standard Bayesian optimisation typically uses multiple randomly generated data points to initialise the algorithm, ensuring a well-fitting surrogate model from the first optimisation step. These initial data points are not selected by maximising the acquisition function but by using space-filling designs, such as Latin hypercube sampling (McKay *et al.*, 1979). These designs assume that outputs for any arbitrary combination of variable values can be observed with experiments. This is not feasible with changing environmental conditions and ENVBO must be initialised differently. Diessner *et al.* (2024) suggest initialisation with just one initial training point, where values for the environmental variables are set to current measurements and the controllable variables are chosen at random. Previous work could not find a significant difference between initialising Bayesian optimisation with many or just a small number of data points (Diessner *et al.*, 2022).

Figure 2 illustrates the general idea of ENVBO's extension to Bayesian optimisation, where plot **A)** shows the true outputs of a 2-dimensional test function considered as an example. The yellow areas indicate high function outputs while the blue areas indicate low function outputs. Assume that x_1 is an environmental variable that cannot be controlled and x_2 is a controllable variable. The main objective is to find the optimal value for the controllable input x_2 for any value for the environmental input x_1. In the example, this is equivalent to learning a function that gives the optimal value of x_2 for any x_1. Assume ENVBO is in the middle of an optimisation run and has already selected 20 data points and observed them from the objective function. The next optimisation step is broken down into two parts in plots **B)** and **C)** of Figure 2. The objective function is represented by a surrogate model fitted to the 20 available data points and a measurement of the environmental variables is taken resulting in $x_1 = -0.5$. ENVBO conditions the optimisation on this measurement by essentially taking a slice through the 2-dimensional variable space and disregarding the rest of

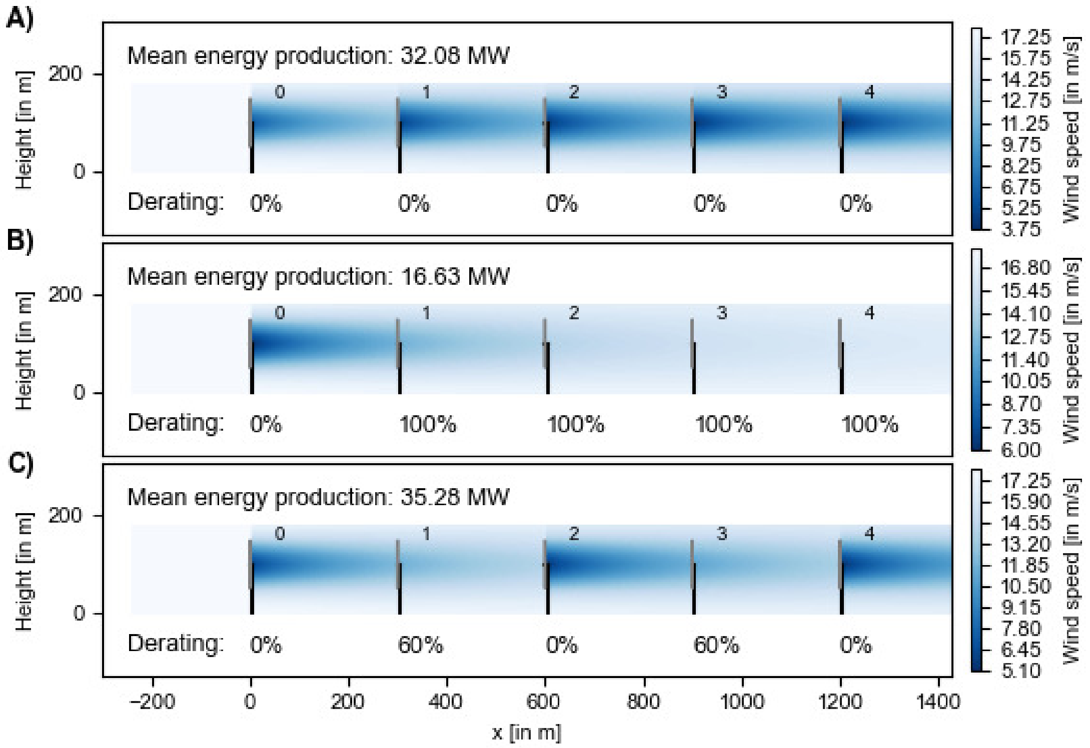

Figure 3: Mean energy production in MW for three different derating strategies with a wind speed of 18 m/s. **A)** All five wind turbines are run at maximum capacity (no derating). **B)** Wind turbine 0 is run at full capacity and wind turbines 1–4 are turned off completely. **C)** Wind turbines 0, 2 and 4 are run at maximum capacity and wind turbines 1 and 3 are reduced to 40%.

the variable space for the current optimisation step. Conditioning on the measurement—i.e., holding the uncontrollable variable fixed at the most recent measurement—reduces the dimensionality of the problem while using all available information. The 1-dimensional problem can now be optimised for x_2 with Bayesian optimisation as shown in plot **C)**. The optimisation step resembles standard Bayesian optimisation given in Figure 1 except that no available data points lie directly on the 1-dimensional slice. Despite this, the surrogate model still offers valuable information about the objective function due to correlation with the environmental variable, showing that combining controllable and environmental variables into a single optimisation strategy maximises the available information of the algorithm.

EXAMPLE APPLICATION

This section illustrates the performance of ENVBO by considering a 5-dimensional wind turbine simulator implemented in PyWake (Pedersen *et al.*, 2023). The use case assumes a row of five wind turbines with the wind blowing directly towards the turbines so that previous wind turbines' wakes affect subsequent wind turbines' energy production. The objective is to maximise the mean energy production (MEP) over the five wind turbines by derating one or multiple wind turbines. The derating—i.e., running wind turbines below their maximum capacity for a given wind speed—reduces the wake of the turbine and increases the potential energy generation of wind turbines affected by the wake downstream. Derating can also extend the lifetime of the wind turbine components (Boersma *et al.*, 2017; Juangarcia *et al.*, 2018; Vernica *et al.*, 2018). Figure 3 shows three different strategies and their resulting mean energy production for a fixed wind speed of 18 m/s. Strategy **A)** runs all five wind turbines at full capacity with no derating and produces 32.08 MW. Strategy **B)** only runs the first wind turbine at full capacity and turns off the four other wind turbines. The wake of wind turbine 0 would affect the energy production potential of wind turbines 1 and 2 as the lower local wind speeds show. This strategy produces 16.63 MW. Strategy **C)** derates wind turbines 1 and 3 by 60% and runs the other three wind turbines at full capacity. This produces 35.28 MW energy and is superior to Strategy **A)**, showing that naively running all wind turbines at full capacity is not necessarily the best approach.

This study uses a fictional wind turbine with a height and diameter of 100 metres for which the effect of derating is computed using 1-dimensional momentum theory as outlined in the documentation of PyWake (Pedersen *et al.*, 2023). The maximum energy generation of a wind turbine is reached at 20 m/s. The five derating levels are controllable variables bounded between 0–100% while the wind direction is fixed at 270 degrees and the wind speed is an environmental variable that changes according to a random walk (Papoulis, 1965). The random walk uses the wind speed of the previous run and adds to it a small value sampled from a uniform distribution $\mathscr{U}[-5,5]$. The maximal change from one iteration to the next is ± 5 m/s and the wind speed is bounded between 6–50 m/s.

ENVBO is benchmarked against two optimisation algorithms—the well-established Nelder-Mead algorithm (Nelder & Mead, 1965) and standard Bayesian optimisation with expected improvement as outlined in this article. While ENVBO is capable of predicting optimal derating combinations for the entire range of wind speeds, Nelder-Mead and standard Bayesian optimisation are not, and

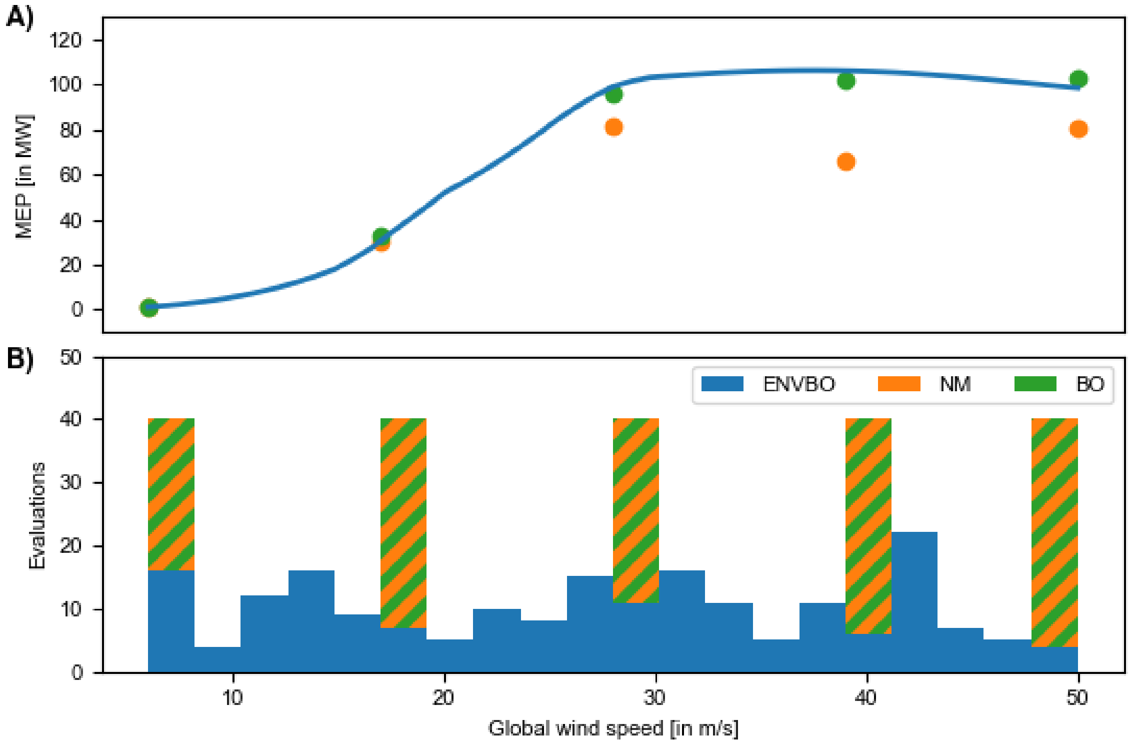

Figure 4: Results of ENVBO against two benchmarks—the Nelder-Mead algorithm and standard Bayesian optimisation. **A)** Mean energy production in MW for different wind speeds. Only ENVBO can predict solutions over the full wind speed range. **B)** Number of evaluations per wind speed.

have to be run for each individual wind speed, keeping it fixed for the full optimisation run. Five different wind speeds—6, 17, 28, 39 and 50 m/s—are taken as samples to enable the comparison of ENVBO to the benchmarks. ENVBO is still run over the entire wind speed range between 6–50 m/s but Nelder-Mead and standard Bayesian optimisation are run separately for these five fixed wind speeds. The results of this comparison are given in Table 1. ENVBO outperforms Nelder-Mead by 17.8%, 2.2%, 21.9%, 60.0% and 21.8% for each wind speed respectively. Compared to standard Bayesian optimisation, ENVBO performs better for wind speeds of 28 and 39 m/s and slightly worse but still comparable for wind speeds 6, 17 and 50 m/s. ENVBO also finds the highest overall mean energy production at 105.99 MW for a wind speed of 39 m/s.

Figure 4 illustrates these results through two plots. Plot **A)** plots the mean energy production in MW for ENVBO, Nelder-Mead (NM) and standard Bayesian optimisation (BO). This shows that ENVBO and standard Bayesian optimisation come to very similar results but Nelder-Mead is clearly worse—particularly for wind speeds 28, 39 and 50 m/s. It also presents the difference in the number of solutions between ENVBO and the benchmarks. ENVBO can predict solutions for any wind speed—here depicted by a line—while the benchmarks only give results for the five specific wind speeds. Plot **B)** connects this performance with the number of function evaluations necessary to get the results. ENVBO was limited to a budget of 200 evaluations and the individual runs of Nelder-Mead and standard Bayesian optimisation were restricted to 40 function evaluations to achieve the same evaluation budget as ENVBO. Plot **B)** shows that the benchmarks have more evaluations available for the five fixed wind speeds. ENVBO's largest number of evaluations falls between

Table 1: Mean energy production of wind turbines optimised with ENVBO, the Nealder-Mead algorithm and standard Bayesian optimisation for the five considered wind speeds.

Wind speed [in m/s]	ENVBO [in MW]	Nelder-Mead [in MW]	BO [in MW]
6	1.19	1.01	1.37
17	30.75	30.10	32.69
28	98.91	81.17	95.90
39	105.99	66.23	101.89
50	98.37	80.78	102.36

wind speeds 41.2–43.4 m/s with 22 evaluations. For the other bins of the histogram, the number of function evaluations is as low as 4 for wind speeds of 8.2–10.4 m/s. Despite having fewer data points available per fixed wind speed—ENVBO does not have a single data point with an exact wind speed of 17, 28 or 39 m/s—ENVBO still finds solutions that outperform or are at least comparable to the benchmarks. For the predictions for specific wind speeds, ENVBO leverages the correlation of the environmental variable and draws from the information of the surrounding wind speeds, requiring much fewer data points to give a solution. This also explains why ENVBO performs more poorly for a wind speed of 50 m/s than standard Bayesian optimisation in this example. ENVBO has very few data points available for that wind speed and, as 50 m/s is the upper bound of the wind speed, can only draw from information from lower wind speeds. While ENVBO performs comparably to the best benchmark, it should be a point of caution that ENVBO might be less accurate towards the boundaries of the environmental variables and extrapolation should be particularly avoided. Overall, ENVBO presents as a sample-efficient and cost-effective optimisation strategy for experiments and simulations with changing environmental variables that performs well compared to well-established benchmarks.

CONCLUSION AND OUTLOOK

Uncontrollable environmental variables can be as influential as controllable variables in many experiments that are conducted in dynamic environments. While environmental factors such as temperature and humidity can affect experimental outputs in a controlled lab setting, they grow even more influential when transitioning towards more realistic experiments where environments cannot be controlled. This study investigates ENVBO—a novel approach for sample-efficient Bayesian optimisation with changing environmental conditions—that takes account of controllable variables and uncontrollable environmental variables in one surrogate model but optimises only the controllable variables conditionally on measurements of the uncontrollable variables. This leverages all available information and provides a promising approach to optimise experiments in dynamic environments with changing environmental conditions while keeping costs low. The approach was applied to a wind farm simulator where the derating of five wind turbines was controlled to maximise the mean energy production for different wind speeds. ENVBO found strategies that outperformed the popular Nelder-Mead algorithm by 2.2–60.0% and were comparable to standard Bayesian optimisation. In contrast to these benchmarks, ENVBO is capable of predicting optimal derating combinations for the entire range of wind speeds. The aim is to apply this approach to physical experiments in the wind tunnel with randomly changing free-stream velocities to maximise drag reduction by actively controlling blowing actuators.

ACKNOWLEDGEMENTS

The work has been supported by the Engineering and Physical Sciences Research Council (EPSRC) under grant numbers EP/T020946/1, EP/T021144/1 and the EPSRC Centre for Doctoral Training in Cloud Computing for Big Data under grant number EP/L015358/1.

REFERENCES

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

Bushnell, D. M. & Hefner, J. N. 1990 *Viscous drag reduction in boundary layers*. American Institute of Aeronautics and Astronautics.

Diessner, M., O'Connor, J., Wynn, A., Laizet, S., Guan, Y., Wilson, K. J. & Whalley, R. D. 2022 Investigating Bayesian optimization for expensive-to-evaluate black box functions: Application in Fluid Dynamics. *Frontiers in Applied Mathematics and Statistics* **8**.

Diessner, M., Wilson, K. J. & Whalley, R. D. 2023 NUBO: A transparent Python package for Bayesian optimisation. *arXiv preprint arXiv:2305.06709* pp. 1–21 [submitted to Journal of Statistical Software].

Diessner, M., Wilson, K. J. & Whalley, R. D. 2024 On the development of a practical Bayesian optimisation algorithm for expensive experiments and simulations with changing environmental conditions. *arXiv preprint arXiv:2402.03006* pp. 1–23 [submitted to Data–Centric Engineering].

Frazier, P. I. 2018 A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811* pp. 1–22.

Hwang, D. 2004 Review of research into the concept of the microblowing technique for turbulent skin friction reduction. *Progress in Aerospace Sciences* **40** (8), 559–575.

Jones, D. R., Schonlau, M. & Welch, W. J. 1998 Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* **13** (4), 455.

Juangarcia, D. A., Eguinoa, I. & Knudsen, T. 2018 Derating a single wind farm turbine for reducing its wake and fatigue. *Journal of Physics: Conference Series* **1037**.

Kornilov, V. I. & Boiko, A. V. 2012 Efficiency of air microblowing through microperforated wall for flat plate drag reduction. *AIAA Journal* **50** (3), 724–732.

Krause, Andreas & Ong, Cheng 2011 Contextual Gaussian process bandit optimization. *Advances in Neural Information Processing Systems* **24**.

Mahfoze, O. A., Moody, A., Wynn, A., Whalley, R. D. & Laizet, S. 2019 Reducing the skin-friction drag of a turbulent boundary-layer flow with low-amplitude wall-normal blowing within a Bayesian optimization framework. *Physical Review Fluids* **4** (9), 094601.

Mallor, F., Semprini-Cesari, G., Mukha, T., Rezaeiravesh, S. & Schlatter, P. 2023 Bayesian optimization of wall-normal blowing and suction-based flow control of a NACA 4412 wing profile. *Flow, Turbulence and Combustion* pp. 1–26.

McKay, M. D., Beckman, R. J. & Conover, W. J. 1979 A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21** (2), 239–245.

Nelder, J. A. & Mead, R. 1965 A simplex method for function minimization. *The Compututer Jorunal* **7**, 308–313.

O'Connor, J., Diessner, M., Wilson, K. J., Whalley, R. D., Wynn, A. & Laizet, S. 2023 Optimisation and analysis of streamwise-varying wall-normal blowing in a turbulent boundary layer. *Flow, Turbulence and Combustion* pp. 1–29.

Papoulis, Athanasios 1965 *Probability, random variables and stochastic processes*. New York, New York: McGraw Hill.

Pedersen, M. M., Forsting, A. M., van der Laan, P., Riva, R., Romàn, L. A. A., Risco, J. C., Friis-Møller, M., Quick, J., Christiansen, J. P. S., Rodrigues, R. V., Olsen, B. T. & Réthoré, P. 2023 PyWake 2.5.0: An open-source wind farm simulation tool.

Rasmussen, C. Edward. & Williams, C. K. I. 2006 *Gaussian processes for machine learning*. The MIT Press Cambridge, MA.

Snoek, J., Larochelle, H. & Adams, R. P. 2012 Practical Bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems* **25**.

Vernica, I., Ma, K. & Blaabjerg, F. 2018 Optimal derating strategy of power electronics converter for maximum wind energy production with lifetime information of power devices. *IEEE Journal of Emerging and Selected Topics in Power Electronics* **6**, 267–276.