

Bayesian Optimisation Against Climate Change: Applications and Benchmarks

Sigrid Passano Hellan¹[0009–0002–6443–0395], Christopher G. Lucas²[0000–0002–6655–8627], and Nigel H. Goddard²[0000–0002–1731–3412]

¹ NORCE Norwegian Research Centre sipa@norceresearch.no
and Bjerknes Centre for Climate Research, Bergen, Norway

² School of Informatics, University of Edinburgh, UK

Abstract. Bayesian optimisation is a powerful method for optimising black-box functions, popular in settings where the true function is expensive to evaluate and no gradient information is available. Bayesian optimisation can improve responses to many optimisation problems within climate change for which simulator models are unavailable or expensive to sample from. While there have been several demonstrations of climate-related applications, there has been no unifying review of applications and benchmarks. We provide such a review here, to encourage the use of Bayesian optimisation for important and well-suited applications. We identify four main application domains: material discovery, wind farm layout, optimal renewable control and environmental monitoring. For each domain we identify a public benchmark or data set that is easy to use and evaluate systems against, while being representative of real-world problems. Due to the lack of a suitable benchmark for environmental monitoring, we propose LAQN-BO, based on air pollution data. Our contributions are: a) summarising Bayesian optimisation applications related to climate change; b) identifying a representative range of benchmarks, providing example code where necessary; and c) introducing a new benchmark, LAQN-BO.

Keywords: Bayesian optimisation · Surrogate models · Sustainability.

1 Introduction

The use of machine learning (ML) to tackle climate change is gaining traction, including ML-wide surveys of relevant applications [57,21]. To facilitate technical innovation and its adoption into practice, it is important to complement these reviews with concise summaries of problems that specific ML frameworks are well-matched to, along with example data sets. Here, we provide a summary, paired with data sets, for Bayesian optimisation. This is important for three reasons: a) it reduces start-up costs and overhead, as the time-consuming task of identifying suitable data sets does not have to be repeated; b) it helps researchers focus on problems that are timely and/or important, and have a structure that is suitable for the framework(s) they are using; and c) it makes it easier to compare

methods, as numerical results can be compared directly. The benchmarks should be similar to real-world applications. That will provide more realistic expectations of performance and show which versions of Bayesian optimisation are most promising, e.g. in terms of underlying models. It will also identify remaining challenges, e.g. in terms of generating useful priors from related data.

Bayesian optimisation (BO), is a machine learning surrogate-based approach to black-box optimisation [29]. It is particularly suited to problems that:

- a. Have a complex unknown structure which can only be efficiently modelled through a surrogate model, e.g. the conductivity of solar panel materials.
- b. Are expensive or slow to evaluate, requiring sample efficiency from trying out fewer bad or uninformative choices, e.g. when synthesising new materials.
- c. Require identifying extrema, either for optimisation, e.g. maximising power generation, or for intervention, e.g. emission monitoring.

Bayesian optimisation consists of fitting a probabilistic surrogate model, typically a Gaussian process [55], to observed data, and using the surrogate model to choose what input value to evaluate next. Then the surrogate model is adjusted, another input value evaluated, and so on. For example, Bayesian optimisation can be used to find the angle of a solar panel that maximises power output, see Fig. 1: (1) model the output as a function of angle based on all available data and/or domain knowledge; (2) select a promising angle to try next, e.g., an angle that is a good candidate for maximising power output or that is maximally informative about the best angle; and (3) try this angle before updating the model and returning to Step 1. An additional benefit of Bayesian optimisation is that it explicitly models its uncertainty, which can be used to determine when to stop optimising [41], or to trade off exploration and exploitation [18].




Fig. 1: Bayesian optimisation to maximise solar panel power generation.

We present a summary of Bayesian optimisation as applied to climate change. We identify the main application domains, emphasise their climate-related impact, and identify beginner-friendly benchmarks for each. We found a benchmark to be lacking for environmental monitoring, so provide LAQN-BO, based on air pollution data [34]. The four use cases are: *material discovery* — accelerating the development of e.g. new solar panels; *wind farm layouts* — choosing where to place individual turbines; *optimal renewable control* — choosing operating

parameters of solar panels or wind turbines; and *environmental monitoring* — choosing where to place sensors. In this paper we mainly focus on Bayesian optimisation. However, there are other closely related methods such as Bayesian experimental design [53,68] and we have opted to include some examples of these methods as well, e.g. [66,36]. [13] is a related survey of Bayesian optimisation and other surrogate-based optimisation methods. Where [13] takes a higher-level approach, focusing on methodology, we instead group work by applications in order to identify data sets and applications for new projects. When selecting benchmarks, we have emphasised ease of adoption and open access.

Bayesian optimisation is not a complete solution, but an important tool when combined with domain knowledge, and can help improve responses to climate change. Currently, Bayesian optimisation is often evaluated on synthetic functions or hyperparameter optimisation (HPO) benchmarks. But performance in real applications is the ultimate test of any ML method, including Bayesian optimisation. HPO is just one kind of problem and synthetic problems often fail to anticipate crucial challenges, e.g. in terms of finding suitable inductive biases [33]. To increase the impact of BO, and its prevalence, we should adopt a broader set of benchmarks. That would not only demonstrate which methods work well in practice, but open up new research problems to solve. Using benchmarks related to climate change and sustainability has the advantage of doing so in an innovative field, increasing the chance of adoption, and with real impact.

2 Material discovery

Material discovery means finding materials with superior combinations of properties [26], e.g. more efficient solar cells. Bayesian optimisation and related methods have been suggested for optimising a wide range of materials, from concrete [60,5] and direct air capture of CO₂ [49] to solar panel glass [31], electrifying transport and chemical industries [7], and for electrocatalysts for CO₂ reduction and H₂ production [66,74,27]. The benefits are in reducing greenhouse gas emissions or increasing their capture; making renewable energy generation or storage more efficient; or running more sectors off renewable energy. If evaluating the material requires synthesising it, there are huge potential gains in time and cost by only doing so on the most promising candidates.

[38] provide three material discovery data sets [63,11,43] related to solar panels, with 3–5 features and 94–178 unique evaluated data points. Using Fig. 1 to illustrate the Bayesian optimisation process, the steps become: (1) model the material performance as a function of ingredient proportions used; (2) select a new combination of ingredient proportions to try; and (3) produce and evaluate the material. We replace (3) with looking up the value in the benchmark.

- **Benchmark:** <https://github.com/PV-Lab/Benchmarking> [38]
- **Features:** Material properties: 3–5 dimensions
- **Data type:** CSV files
- **Sampling:** Manufacture material. In benchmark replaced by table look-up

- **Objective:** Material performance: conductivity, absorbance or stability
- **Impact:** Better solar panels; more renewables and less climate gases

3 Wind farm layout

Wind turbines are typically grouped together in wind farms, with several turbines in relatively close proximity, to reduce installation costs and environmental impacts. Before construction, the locations of individual wind turbines must be planned. Wind turbines work by extracting kinetic energy from the air, so the wind is weaker and more turbulent after flowing past a turbine. This leads to less power generation [50] and greater dynamic loads and fatigue on the downwind turbines [20]. This is illustrated in Fig. 2, where the two turbines need to be placed within the green area and the wind is shown by arrows. Determining the optimal wind turbine layout is a difficult but important optimisation problem, since more renewable electricity can be generated without requiring more turbines to be installed. The efficiency of Bayesian optimisation is important for this problem, since the type of simulations used can take 15 seconds to run on one CPU even for five wind turbines [14], and the simulations must be run many times to compare layouts. Bayesian optimisation can be applied to the problem by replacing the steps in Fig. 1 with (1) modelling the collective power output as a function of the wind turbine placements; (2) selecting a new combination of locations to evaluate; and (3) running the power output simulation.




Fig. 2: Placement of wind turbines within limits shown in green. The wind is shown by the blue arrows: wider arrows indicate stronger wind and more undulating arrows indicate more turbulence. In the layout on the left one turbine is directly downwind of the other, resulting in weaker and more turbulent wind. To the right, both turbines get unhindered wind. In the benchmark five turbines are placed simultaneously, here we illustrate the problem with two turbines.

[14]’s set of benchmarking problems and baselines include a wind farm layout problem, *windwake*, based on FLORIS [48] simulations. The benchmark assumes a fixed number of wind turbines; this is relaxed in [17], where the two objectives of cost and power generation are jointly optimised, using simulations based on [51]. In related pieces of work, [44] build wind maps for later layout planning, and [65] plan the layout of bifacial solar panel arrays. We use the benchmark in [14] as it has been prepared for further use.

- **Benchmark:** <https://github.com/AlgTUDelft/ExpensiveOptimBenchmark>
- **Features:** Spatial locations: 10 dimensions. 2 for each of 5 turbines [14]
- **Data type:** Simulation
- **Sampling:** Simulate layout
- **Objective:** Energy production
- **Impact:** More renewable generation, which reduces climate gas emissions

4 Optimal renewable control

Having planned the layout for a renewable power farm, the next question is how best to operate it [20]. For wind turbines the blade angles [19,50,71], and induction factor [50] need to be set, and for airborne systems the altitude [54]. For solar panels the voltage level applied [1,40] and the angle of the panel [3] can be adjusted. By adjusting these to maximise energy generation, we minimise emissions for a given electricity demand. For wind farms, the optimisation should occur jointly, as the available wind energy for a downstream turbine is impacted by the operating settings of the upstream turbine [50]. With Bayesian optimisation the response of the non-linear problem can be learnt through data collection. Extensive work on wind turbine control has been done using Bayesian Ascent, a version of BO which limits changes in inputs between iterations, see e.g. [50]. Relatedly, [46] tune the parameters of traditional control methods, and [47] reduce blade fatigue. [6] use a related method for wind turbine control.




Fig. 3: Adjusting solar panels to optimise power generation as the sun moves across the sky. At a simple level, the problem of adjusting angles can be understood as maximising the amount of sunlight hitting the panels, which means the optimal angle varies throughout the day. The benchmark takes the challenge a step further by not only optimising the amount of *direct* sunlight, but also diffusive and reflected sunlight.

[3] provide a benchmark for solar panel optimisation. The goal is to maximise energy yield by adjusting solar panel angles, see Fig. 3. They evaluate contextual bandits and reinforcement learning methods, but the problem is also suitable for BO: we provide a basic example³ at https://github.com/sighellan/solar_panels_rl. Fig. 1 shows how BO is applied to the problem. We model power output as a function of panel angles; decide on a new angle to try; and run the simulation to evaluate the new angle. The challenge is not tracking the sun; its location is predictable and given as an input. Instead, the idea is to adapt to other changes such as cloud cover. However, the angle of many solar panels cannot be adjusted. Wind turbines have more control settings, interact in less predictable ways, and have seen more applications of BO. We chose the benchmark in [3] because it is well-documented and publicly available, but real-world benefits are more likely to come from applications to wind turbines.

- **Benchmark:** https://github.com/david-abel/solar_panels_rl [3]
- **Features:** Panel direction: 2 dimensions
- **Data type:** Simulation
- **Sampling:** Simulate collected sunlight
- **Objective:** Collected energy
- **Impact:** More renewable generation, which reduces climate gas emissions

5 Environmental monitoring

An important problem within environmental monitoring is that of choosing where to place sensors. This is difficult since the planning must be done without knowing the usefulness of each site, which depends on future data. BO and related methods are useful tools in this setting, as they build probabilistic models of the environmental characteristic of interest, e.g. air pollution, which they use to make efficient choices. If we are mainly interested in the worst-hit locations, e.g. for targeting interventions or evaluating legal compliance, BO lets us do that efficiently. To build accurate environmental maps for an entire area we can instead use Bayesian experimental design, e.g. if we want to know the water quality throughout a lake. Monitoring is important for tracking the impacts of climate change, including on oceans [59], for locating gas leakages [8,28,36], and the related task of air pollution monitoring [4,45,42,61,32,33,37]. [58] provide a benchmark for path-planning based on water quality monitoring, also used by [25], but the benchmark uses synthetic data.

The LAQN-BO benchmark: We present a new climate change related benchmark, using air pollution data from the London Air Quality Network (LAQN) [34], see examples in Fig. 4.. Air pollution is related to climate change in that the sources and hence required interventions are often interlinked, e.g.

³ The benchmark relies on the `simple_rl` Python package [2]; we used version 0.784. For our example Bayesian optimisation implementation we additionally relied on BoTorch [10]. The implementation is not meant to be competitive with the existing methods for the benchmark, but rather to demonstrate how BO can be applied.

car exhaust. Also, environmental changes are often sped up by climate change, making their monitoring more important. The optimisation problems in the LAQN-BO benchmark are similar to those in [33]. But while that focuses on methodology, we here focus on the problems themselves. We provide simple scripts and step-by-step instructions for generating the problems and put them in the context of other climate-related BO problems and benchmarks.

The optimisation objective is the NO_2 concentration, i.e. to find the location of maximum pollution from the set of available locations. This is useful as higher pollution concentrations have greater impacts on human health. Using Fig. 1 as a basis, the steps become: (1) model the NO_2 concentration as a function of spatial location; (2) determine a new location to evaluate; and (3) measure the pollution in the new location. We use LAQN-BO to replace (3) with a historical measurement. We construct training and test sets from the 2015 and 2016 data, respectively. Each problem corresponds to a single day, so multiple days give multiple problems. As in [33], we filter out days when fewer than 40 stations collected measurements, and only use the ‘Roadside’ sensors, resulting in 214 training problems and 365 test problems. We log-transformed the data, and standardised using the mean and standard deviation from the training set.

The LAQN-BO benchmark has several advantages. Firstly, it uses real data. Secondly, we present a *set* of related problems, so that methods can be trained on the training set, and evaluated on the test set. Finally, the provided scripts can be easily adapted to generate new problems using other pollutants or years of the extensive LAQN data. None of the other benchmarks combine all of these advantages. A disadvantage is that data are not available everywhere, only at locations with existing sensors, limiting the available search space. This corresponds to having to choose from a sparse set of locations. For instance, if we are mounting sensors to lampposts, we can only choose between existing lampposts.

- **Benchmark:** <https://github.com/sighellan/LAQN-BO>
- **Features:** Spatial coordinates: 2 dimensions
- **Data type:** Python class
- **Sampling:** Place pollution sensor. In benchmark replaced by table look-up
- **Objective:** NO_2 concentration
- **Impact:** Better monitoring of air pollution, enabling targeted interventions

6 Benchmark comparison

The four benchmarks identified in our survey have different characteristics, as is summarised in Table 1. They represent different key challenges for Bayesian optimisation deployments, and provide breadth for evaluating different types of Bayesian optimisation methods. The key challenge of the materials benchmark is its small size, meaning there is little opportunity to fine-tune the methods. For LAQN-BO it is that of learning from the training data, as the problems have as little as 40 evaluations each, requiring very sample-efficient learning. The wind farm layout benchmark also requires efficient use of data. But in contrast




Fig. 4: Example LAQN-BO problems from the training set (2015, top) and test set (2016, bottom). Not all sensors are available on every day. The overall pollution level varies, compare the top left and right plots. The locations of maxima also vary, compare the right top and bottom plots. Most of the sensors are clustered together in central London.

to LAQN-BO, it does not provide training data for transfer learning. The key challenge for the renewable control benchmark is dealing with large numbers of samples, as the panel direction is optimised regularly over many days. It is therefore a good candidate for testing Bayesian optimisation built on scalable Gaussian processes [39]. It also requires keeping track of context, e.g. time, as the sun moves across the sky. An additional contrast to the other benchmarks is that the performance throughout the optimisation matters, as it directly impacts the amount of electricity generated.

Table 1: Comparison of benchmarks.

Benchmark	Type	Key challenge	Dimensionality	# problems
Materials [38]	Observations	Small size	3/5	3
Wind farm layout [14]	Simulation	Sample efficiency	10	1
Renewable control [3]	Simulation	Scaling, context	2	10
Environmental monitoring	Observations	Transfer learning	2	214+365

The materials benchmark is the easiest to start using, as the data are provided in CSV files, followed by that for environmental monitoring, LAQN-BO, which provides a simple Python interface to the data. The benchmarks for wind

farm layout and renewable control require more setup, as methods need to be interfaced with the simulators. Their advantage is that they are easier to extend, and can be evaluated for more input values. LAQN-BO can be easily extended to more problems using the LAQN data [34], but extensions to other base data sets would require more work. The materials benchmark is the hardest to extend, as it relies on results from synthesising and evaluating materials.

7 Conclusion and other applications

There are many more climate-related applications: learning energy consumptions of household appliances [35]; optimising charging protocols for electric vehicles [9]; scheduling smart appliances to smooth out demand curves [64]; improving simulations [56] and experiments [30] for nuclear fusion; reducing vibrations for floating wind turbines [73]; reducing the impact of livestock diseases made more prevalent by global warming [62]; and tuning parameters for heating and cooling system and building models to reduce energy demands [72,15,23]. Bayesian optimisation also has applications in improving climate models, by targeting informative training data [69], and in polar climate science [12]. And it has been used for parameter tuning of a global land surface model [22], a landslide model [52] and forecasting models for electricity generation and demand [67].

To encourage more Bayesian optimisation practitioners to work on climate change applications, it is important to establish more benchmarks, standard data sets or simulators, as we do here with the LAQN-BO benchmark. More challenging benchmarks should also be established, e.g. for constrained or multi-objective optimisation, to encompass more fully the complexity of real problems. For example: some of the area might be unsuitable for building wind turbines; and when building generators there are multiple competing objectives [70,24]. Additionally, there are constraints for controlling the generators, due to increased energy costs and fatigue when changing settings. The energy cost of changing angles is included in the renewable control benchmark [3], so Bayesian Ascent [50] should be tested, as it limits the difference between iterations. As a further extension, the layout and control of wind farms can be optimised jointly [16].

In the future, all these benchmarks should be brought together and their interfaces aligned, to create a standardised, climate-themed benchmark suite. By identifying four benchmarks representing the main application domains, introducing a new benchmark, and highlighting the potential for Bayesian optimisation, we take an important step in facilitating future research.

Acknowledgments. We thank Linus Ericsson, Simon Valentin and Michael Gutmann for helpful suggestions. This work was supported by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1) and the University of Edinburgh.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdelrahman, H., et al.: Bayesian optimization for maximum power point tracking in photovoltaic power plants. *European Control Conference* (2016)
2. Abel, D.: `simple_rl`: Reproducible reinforcement learning in python. *ICLR Workshop on Reproducibility in Machine Learning* (2019)
3. Abel, D., et al.: Bandit-based solar panel control. In: *AAAI* (2018)
4. Ainslie, B., Reuten, C., Steyn, D.G., Le, N.D., Zidek, J.V.: Application of an entropy-based Bayesian optimization technique to the redesign of an existing monitoring network for single air pollutants. *Journal of Environmental Management* (2009)
5. Ament, S., Witte, A.C., Garg, N., Kusuma, J.: Sustainable concrete via Bayesian optimization. *NeurIPS 2023 RealML Workshop* (2023)
6. Andersson, L.E., Imsland, L.: Real-time optimization of wind farms using modifier adaptation and machine learning. *Wind Energy Science* (2020)
7. Annevelink, E., et al.: Automat: Automated materials discovery for electrochemical systems. *MRS Bulletin* (2022)
8. Asenov, M., et al.: Active Localization of Gas Leaks Using Fluid Simulation. *IEEE Robotics and Automation Letters* (2019)
9. Attia, P.M., et al.: Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature* (2020)
10. Balandat, M., et al.: BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. In: *NeurIPS* (2020)
11. Bash, D., et al.: Multi-fidelity high-throughput optimization of elect. conductivity in P3HT-CNT composites. *Advanced Functional Materials* (2021)
12. Bente, K., et al.: Probabilistic ML in polar earth and climate science (2022)
13. Bliet, L.: A survey on sustainable surrogate-based opt. *Sustainability* (2022)
14. Bliet, L., et al.: Benchmarking surrogate-based optimisation algorithms on expensive black-box functions. *Applied Soft Computing* (2023)
15. Chakrabarty, A., et al.: Scalable Bayesian optimization for model calibration. *Energy and Buildings* (2021)
16. Chen, K., Lin, J., Qiu, Y., Liu, F., Song, Y.: Joint optimization of wind farm layout considering optimal control. *Renewable Energy* (2022)
17. Chugh, T., Ymeraj, E.: Wind farm layout optimisation using set based multi-objective Bayesian optimisation. In: *GECCO* (2022)
18. De Ath, G., et al.: Greed is good: Exploration and exploitation trade-offs in Bayesian opt. *Evolutionary Learning and Optimization ACM* (2021)
19. Doekemeijer, B.M., et al.: Model-based closed-loop wind farm control for power maximization using Bayesian optimization. In: *IEEE CCTA* (2019)
20. Dong, H., Xie, J., Zhao, X.: Wind farm control technologies: From classical control to reinforcement learning. *Progress in Energy* (2022)
21. Donti, P.L., Kolter, J.Z.: Machine learning for sustainable energy systems. *Annual Review of Environment and Resources* (2021)
22. Druel, A., et al.: Towards a more detailed representation of high-latitude vegetation in the global land surface model ORCHIDEE (ORC-HL-VEGv1. 0). *Geoscientific Model Development* (2017)
23. Fiducioso, M., et al.: Safe contextual Bayesian optimization for sustainable room temperature PID control tuning. In: *IJCAI* (2019)
24. Flecker, A.S., et al.: Reducing adverse impacts of Amazon hydropower expansion. *Science* (2022)

25. Folch, J.P., et al.: SnAKe: Bayesian optimization with pathwise exploration. *NeurIPS* (2022)
26. Frazier, P.I., Wang, J.: Bayesian optimization for materials design. In: *Information science for materials discovery and design*. Springer (2016)
27. Frey, D., et al.: Bayesian optimization of electrochemical devices for electrons-to-molecules conversions. *ChemRxiv* (2022)
28. Gao, T., Bai, X.: Bayesian opt.-based three-dimensional, time-varying env. monitoring using an UAV. *Journal of Intelligent & Robotic Systems* (2022)
29. Garnett, R.: *Bayesian optimization*. Cambridge University Press (2023)
30. Gundecha, V., et al.: Meta-learned Bayesian optimization for energy yield in inertial confinement fusion. *NeurIPS CCAI Workshop* (2024)
31. Haghanifar, S., Cheng, B., McCourt, M., Leu, P.: Using Bayesian optimization to improve solar panel performance by developing antireflective, superomniphobic glass. *ICML CCAI Workshop* (2019)
32. Hellan, S.P., Lucas, C.G., Goddard, N.H.: Optimising placement of pollution sensors in windy environments. *NeurIPS AI4Earth Workshop* (2020)
33. Hellan, S.P., Lucas, C.G., Goddard, N.H.: Bayesian optimisation for active monitoring of air pollution. *AAAI* (2022)
34. Imperial College London, E.R.G.: London Air Quality Network (1993 ongoing)
35. Jia, Y., et al.: Active collaborative sensing for energy breakdown. In: *ACM Information and Knowledge Management* (2019)
36. Kleinegesse, S., Gutmann, M.U.: Bayesian experimental design for implicit models by mutual information neural estimation. In: *ICML* (2020)
37. Leu, F.Y., Ho, J.S.: Air pollution source identification by using neural network with Bayesian optimization. *IMIS* (2020)
38. Liang, Q., et al.: Benchmarking the performance of Bayesian opt. across multiple exp. materials science domains. *npj Comp. Materials* (2021)
39. Liu, H., et al.: When Gaussian process meets big data: A review of scalable GPs. *IEEE neural networks and learning systems* (2020)
40. Lyden, S., Olding, W., Darbari, Z.: Bayesian optimization for maximum power point tracking in photovoltaic systems. In: *PESGM. IEEE* (2018)
41. Makarova, A., et al.: Automatic termination for hyperparameter optimization. *AutoML* (2022)
42. Marchant, R., Ramos, F.: Bayesian optimisation for Intelligent Environmental Monitoring. In: *Intelligent Robots and Systems. IEEE* (2012)
43. Mekki-Berrada, F., et al.: Two-step machine learning enables optimized nanoparticle synthesis. *npj Computational Materials* (2021)
44. Mern, J., Yildiz, A., Sunberg, Z., Mukerji, T., Kochenderfer, M.J.: Bayesian optimized Monte Carlo planning. In: *AAAI* (2021)
45. Morere, P., et al.: Sequential Bayesian optimization as a POMDP for environment monitoring with UAVs. In: *ICRA. IEEE* (2017)
46. Moustakis, N., et al.: A practical Bayesian optimization approach for the optimal estimation of the rotor effective wind speed. In: *ACC. IEEE* (2019)
47. Mulders, S., et al.: Efficient tuning of individual pitch control. In: *Journal of Physics: Conference Series. IOP Publishing* (2020)
48. NREL: FLORIS. version 2.1.1 (2020), <https://github.com/NREL/floris>
49. Ortiz-Montalvo, D.L., et al.: Machine learning for climate change: Guiding discovery of sorbent materials for direct air capture of CO₂. In: *ICML CCAI Workshop* (2021)

50. Park, J., Law, K.H.: Bayesian ascent: A data-driven optimization scheme for real-time control with application to wind farm power maximization. *IEEE Transactions on Control Systems Technology* (2016)
51. Pedersen, M., et al.: Dtuwindenergy/pywake: Pywake. Zenodo [code] (2019)
52. Pradhan, B., et al.: A meta-learning approach of optimisation for spatial prediction of landslides. *Remote Sensing* (2021)
53. Rainforth, T., et al.: Modern Bayesian exp. design. *Statistical Science* (2024)
54. Ramesh, S.S., et al.: Movement penalized Bayesian optimization with application to wind energy systems. *NeurIPS* (2022)
55. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for ML*. MIT (2006)
56. Rodriguez-Fernandez, P., et al.: Enhancing predictive capabilities in fusion burning plasmas through surrogate-based opt. [...]. *Nuclear Fusion* (2024)
57. Rolnick, D., et al.: Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)* (2022)
58. Samaniego, F.P., et al.: A Bayesian opt. approach for water resources monitoring through an autonomous surface vehicle. *IEEE Access* (2021)
59. Sanchez-Pi, N., et al.: Artificial intelligence, machine learning and modeling for understanding the oceans and climate change. In: *NeurIPS CCAI Workshop* (2020)
60. Severson, K., et al.: Amortized inference of Gaussian process hyperparameters for improved concrete strength trajectory prediction. *NeurIPS CCAI workshop* (2021)
61. Singh, A., et al.: Modeling and decision making in spatio-temporal processes for environmental surveillance. In: *IEEE Robotics and Automation* (2010)
62. Spooner, T., et al.: Bayesian optimisation of restriction zones for bluetongue control. *Scientific reports* (2020)
63. Sun, S., et al.: A data fusion approach to optimize compositional stability of halide perovskites. *Matter* (2021)
64. Tabakhi, A., et al.: The smart appliance scheduling problem: A Bayesian optimization approach. *Principles and Practice of Multi-Agent Systems* (2020)
65. Tillmann, P., et al.: Minimising the levelised cost of electricity for bifacial solar panel arrays using Bayesian opt. *Sustainable Energy & Fuels* (2020)
66. Tran, K., Ulissi, Z.W.: Active learning across intermetallics to guide discovery of electrocatalysts for CO₂ reduction and H₂ evolution. *Nature Catalysis* (2018)
67. Trivedi, R., Khadem, S.: Peak demand management and schedule optimisation for energy storage through the machine learning approaches. *EUROCON IEEE Smart Technologies* (2021)
68. Valentin, S., et al.: Designing optimal behavioral exp. using ML. *Elife* (2024)
69. Watson-Parris, D.: Machine learning for weather and climate are worlds apart. *Philosophical Transactions of the Royal Society A* (2021)
70. Wu, X., et al.: Efficiently approximating the pareto frontier: hydropower dam placement in the Amazon basin. In: *AAAI* (2018)
71. Yang, S., et al.: Cooperative yaw control of wind farm using a double-layer machine learning framework. *Renewable Energy* (2022)
72. Zhan, S., et al.: Calibrating building simulation models using multi-source datasets and meta-learned Bayesian opt. *Energy and Buildings* (2022)
73. Zhang, X., et al.: Vibration suppression of multi-component floating structures via passive TMDs and Bayesian ascent. *Ocean Engineering* (2022)
74. Zhong, M., et al.: Accelerated discovery of CO₂ electrocatalysts using active machine learning. *Nature* (2020)