Context-Dependent Bayesian Optimization in Real-Time Optimal Control: A Case Study in Airborne Wind Energy Systems

Ali Baheri
Department of Mechanical Engineering
University of North Carolina at Charlotte
Charlotte, NC 28223
akhayatb@uncc.edu

Chris Vermillion
Department of Mechanical Engineering
University of North Carolina at Charlotte
Charlotte, NC 28223
cvermill@uncc.edu

Abstract

We present a framework in which Bayesian Optimization is used for real-time optimal control. In particular, Bayesian Optimization is applied to the real-time altitude optimization of an airborne wind energy (AWE) system, for the purpose of maximizing net energy production.

1 Introduction

Airborne Wind Energy (AWE) systems are a new paradigm for wind turbines in which the structural elements of conventional wind turbines are replaced with tethers and a lifting body to harvest wind power from significantly increased altitudes (typically up to 600m). At those altitudes, winds are stronger and more consistent than ground-level winds. Besides being able to operate at much higher altitudes than traditional turbines, AWE systems also provide additional control degrees of freedom that allow the systems to *adjust* their operating altitudes and intentionally induce *crosswind motions* to enhance power output. A fundamental challenge lies in choosing what altitude do we go to next, particularly in a partially observable wind speed vs. altitude (wind shear) profile.

For the altitude optimization problem at hand, it is desirable to employ a control system that can learn the statistical properties of the wind shear profile online, thereby alleviating the need for offline data collection prior to running the control system. One of the most studied problems in the machine learning community is the design of optimization algorithms for real-world applications using scarce data. In existing literature, this problem has been studied in the context of sequential decision-making problems aimed at learning the behavior of an objective function (called exploration) while simultaneously trying to optimize performance (called exploitation). As an efficient and systematic approach for balancing exploration and exploration in a partially observable environment, Bayesian Optimization has been applied to various real-world problems [1, 2, 3, 4]. In general, Bayesian Optimization aims to find the global optimum of an unknown, expensive-to-evaluate, and black-box function within only a few evaluations.

1.1 Energy Generation Model

In this section, we formalize the objective function for the altitude optimization problem. An appropriate objective function should account for the energy generated by the turbine and the energy lost by adjusting and maintaining altitude (even when maintaining altitude, control energy must be expended to make small adjustments to the tethers in order to reject typical levels of turbulence). In order to account for each of these factors, the instantaneous net power generation is expressed as [6]:

$$P_{\text{total}} = c_1 \min\big(V_{\text{wind}}, V_{\text{rated}}\big)^3 - c_2 V_{\text{wind}}^2 + P_z\big(V_{\text{wind}}, \dot{z}\big), \tag{1}$$

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

Figure 1: Altaeros Buoyant Airborne Turbine (BAT), Image Credit: [5]

In (1), z is the operating altitude, $V_{\text{wind}}(z)$ is the wind speed at the operating altitude, V_{rated} is the turbine's rated wind speed, P_z is the instantaneous power required for or regenerated by adjusting altitude, and $c_{1,2}$ are constant parameters that depend upon the turbine's power curve and flight control system, respectively. We will use a simplified version of (1) that lumps motoring and regenerated power into a single term that penalizes instantaneous altitude adjustment.

$$P = c_1 \min\Big(V_{\text{wind}}(z), V_{\text{rated}}\Big)^3 - c_2 V_{\text{wind}}^2 - \bar{c}_3 V_{\text{wind}}^2 \mid \dot{z} \mid, \tag{2}$$

Ultimately, equation (2) serves as our objective function, which should be maximized.

1.2 Bayesian Optimization

Suppose that we want to maximize an unknown, black-box, and expensive-to-evaluate objective function. Due to the cost associated with evaluating this function, it is crucial to select the location of each new evaluation deliberately. Bayesian Optimization represents a promising candidate strategy for this altitude optimization, as Bayesian Optimization focuses specifically on converging to the *global optimum* of an unknown function within *a minimal number of evaluations* on the real system. The major assumption here is that function evaluation is expensive, whereas computational resources are cheap. This matches the altitude optimization problem at hand, where each evaluation of the objective function requires the physical AWE system to relocate itself to a new altitude, necessitating substantial time and energy expenditure. Broadly speaking, there are two main components to Bayesian Optimization. First, we need an appropriate model of the objective function. Second, we need to choose an *acquisition function*, which characterizes the fitness of each candidate value of the next operating altitude. The maximization of the expected value of this acquisition function determines the next point (altitude) to operate the system at.

1.2.1 Learning Phase: Using Gaussian Processes (GPs)

In general, a GP is fully specified by its mean function, $\mu(z)$, which is assumed to be zero without loss of generality [7] (a coordinate shift can be used to generate a zero mean objective function), and covariance function, $k(z, z')$:

$$P(z) \sim \mathcal{GP}\Big(\mu(z), k(z, z')\Big), \tag{3}$$

In the context of altitude optimization problem, the GP framework is used to predict the generated power, $P(z)$, at candidate altitude, z_c, based on a set of t past observations, $\mathcal{D}_{1:t} = \{z_{1:t}, P(z_{1:t})\}$.

The power at candidate altitude z_c is modeled to follow a multivariate Gaussian distribution [7]:

$$\begin{bmatrix} y_t \\ P(z_c) \end{bmatrix} \sim \mathcal{N}\Big(0, \begin{bmatrix} K_t + \sigma_\epsilon^2 I_t & k_t^T \\ k_t & k(z_c, z_c) \end{bmatrix}\Big), \tag{4}$$

where $y_t = \{P(z_1), \cdots, P(z_t)\}$ is the vector of observed function values. Moreover, the vector $k_t(z) = [k(z_c, z_1), \cdots, k(z_c, z_t)]$ encodes the covariances between the candidate altitude, z_c, and the past data points, $z^{1:t}$. The past-data covariance matrix, with entries $[K_t]_{(i,j)} = k(z_i, z_j)$ for $i, j \in \{1, \cdots, t\}$, characterizes the covariances between pairs of past data points. The identity matrix is represented by I_t and σ_ϵ represents the noise variance [7].

A common choice for the kernel function is a Squared Exponential (SE) kernel, which takes the form:

$$k(z_i, z_j) = \sigma_0^2 \exp\Big(-\frac{1}{2}(z_i - z_j)^T \Lambda^{-2}(z_i - z_j)\Big). \tag{5}$$

Kernel hyper-parameters are identified by maximizing the marginal log-likelihood of the existing observed data, $\mathcal{D}$ [7]:

$$\theta^* = \arg\max_{\theta} \log p(y_t \mid z^{1:t}, \theta) \tag{6}$$

Once the hyper-parameters have been identified, the predictive mean and variance at z_c, conditioned on these past observations, can be expressed as:

$$\mu_t(z_c \mid \mathcal{D}) = k_t(z_c)\Big(K_t + I_t\sigma_\epsilon^2\Big)^{-1} y_t^T, \tag{7}$$

$$\sigma_t^2(z_c \mid \mathcal{D}) = k(z_c, z_c) - k_t(z_c)\Big(K_t + I_t\sigma_\epsilon^2\Big)^{-1} k_t^T(z_c), \tag{8}$$

1.2.2 Optimization Phase

The ultimate goal of Bayesian Optimization is to determine the next point based on past observations. Among several choices of acquisition functions, we use an acquisition function belonging to the improvement-based family [8]. More precisely, the next operating point is selected as the one that maximizes the so-called expected improvement:

$$z_{t+1} = \arg\max_{z} \mathbb{E}\Big(\max\{0, P_{t+1}(z) - P(z)^{max}\} \mid \mathcal{D}_{1:t}\Big) \tag{9}$$

where $\max\{0, P_{t+1}(z) - P(z)^{max})\}$ represents the *improvement* toward the best value of the objective function so far, $P(z)^{max}$. This quantity is referred to as *expected improvement*.

1.3 Context-Dependent Optimization

The time-varying nature of our objective function (the wind speed at a given altitude does not remain constant over time) compels us to modify the conventional Bayesian Optimization method to account for *spatiotemporal* variation of the wind shear profile. One way to cope with the dynamic nature of the objective function is through contextual GP [9]. More precisely, one of the benefits of modeling the objective function using GPs is that we can characterize the statistical properties of the system output based on environmental variables, called *context*. To deal with time-varying nature of our objective function, we consider *time* as the context.

Employing contextual GP is very straightforward. In Section 1.2.1, the kernel of the GP is defined only in terms of altitude. To account for time-dependency, we modify the kernel structure such that it consists of an addition of two kernels, one over altitude, k_z, and one over time, k_t:

$$k\Big((z, t), (z', t')\Big) = k_z(z, z') + k_t(t, t'), \tag{10}$$

Context-Dependent Bayesian Optimization is summarized in Algorithm 1. First, at the very first two time steps, the procedure is initialized by two previously evaluated altitudes and the corresponding function values (line 2). Then, at each time step, the context (time) is fixed at the current instance (line 4). Next, the composite kernel is constructed based upon Eq. (10) (line 5). The GP model and its predictive mean and variance are computed (line 6-7). Then, the acquisition function value is computed and optimization is done only over the altitude space (line 8-9). Finally, the suggested candidate (altitude) and the current time are augmented to the data (line 10), and the process is restarted.

Algorithm 1 Context-Dependent Bayesian Optimization (CDBO)

procedure ALTITUDE OPTIMIZATION WITH CDBO
  $\mathcal{D} \leftarrow$ *Initialize*: $\{(z_{1:2}, t_{1:2}), P(z_{1:2})\}$
  for each time step **do**
    Set the context into the current time
    Construct a composite kernel (Eq. (10))
    Train a GP model from $\mathcal{D}$
    Compute predictive mean and variance of GP
    Compute acquisition function
    Find z^* that optimizes acquisition function
    Append $\{(z^*, t_{\textit{current time}}), P(z^*)\}$ to $\mathcal{D}$

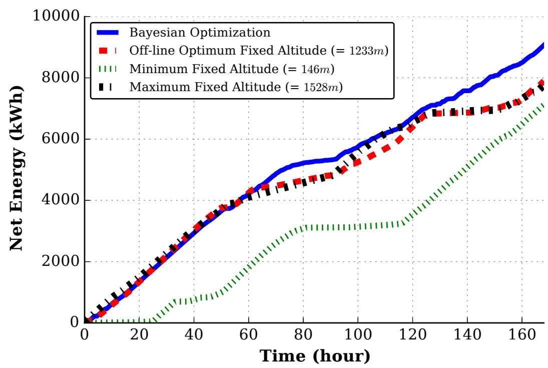

Figure 2: Comparison of net energy production over the course of one week. Context-Dependent Bayesian Optimization outperforms off-line optimum fixed altitude, minimum fixed altitude, and maximum fixed altitude scenarios.

2 RESULTS

We evaluated the proposed context-dependent Bayesian Optimization (CDBO) algorithm by simulating the model based on available data, acquired by a wind profiler in Cape Henlopen, DE [10]. Here, we compare the system performance under four different control scenarios:

- **Scenario 1: Off-line Optimum Fixed Altitude:** The first scenario involves flying the AWE at an *optimum fixed altitude*. We identify this optimal fixed altitude by off-line calculation of energy generation at different altitudes, using given wind data.
- **Scenario 2: Minimum Fixed Altitude:** The second scenario is representative of conventional towered wind turbines. For this scenario, we simulate the performance of the system operating at the lowest available altitude ($146m$), which is comparable to the hub height of the world's tallest towered wind turbines.
- **Scenario 3: Maximum Fixed Altitude:** The third scenario involves flying the AWE at the maximum tabulated fixed altitude, $1528m$.
- **Scenario 4: Bayesian Optimization:** The fourth scenario employs Bayesian Optimization, as outlined in Algorithm 1.

Fig. 2 shows the net energy production across all scenarios. One can conclude from this figure that the CDBO algorithm results in superior net energy production when compared to other scenarios.

3 CONCLUSION

We presented a data-driven approach for optimizing the altitude of an AWE system to maximize the system's net power output. In implementing this altitude optimization approach, we extended conventional Bayesian Optimization to capture time-dependent patterns of our objective function.

References

[1] Ali Baheri, Shamir Bin-Karim, Alireza Bafandeh, and Christopher Vermillion. Real-time control using bayesian optimization: A case study in airborne wind energy systems. *Control Engineering Practice*, 69:131–140, 2017.

[2] Roman Garnett, Michael A Osborne, and Stephen J Roberts. Bayesian optimization for sensor set selection. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, Stockholm, Sweden, 2010.

[3] Ali Baheri, Joe Deese, and Chris Vermillion. Combined plant and controller design using Bayesian optimization: A case study in airborne wind energy systems. In *ASME 2017 Dynamic Systems and Control Conference*, Tysons Corner, US, 2017.

[4] Hany Abdelrahman, Felix Berkenkamp, Jan Poland, and Andreas Krause. Bayesian optimization for maximum power point tracking in photovoltaic power plants. In *Proc. European Control Conference*, Aalborg, DK, 2016.

[5] Altaeros Energies, inc., http://www.altaerosenergies.com.

[6] Alireza Bafandeh and Chris Vermillion. Real-time altitude optimization of airborne wind energy systems using Lyapunov-based switched extremum seeking control. In *American Control Conference*, Boston, US, 2016.

[7] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Number ISBN 0-262-18253-X. The MIT Press, 2006.

[8] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

[9] Andreas Krause and Cheng S Ong. Contextual gaussian process bandit optimization. In *Advances in Neural Information Processing Systems*, pages 2447–2455, 2011.

[10] C. Archer. Wind profiler at Cape Henlopen. Lewes, Delaware, http://www.ceoe.udel.edu/our-people/profiles/carcher/fsmw. Technical report, University of Delaware, 2014.