Control Engineering Practice 69 (2017) 131–140

Contents lists available at ScienceDirect

Control Engineering Practice

journal homepage: www.elsevier.com/locate/conengprac

Real-time control using Bayesian optimization: A case study in airborne wind energy systems☆

Ali Baheri *, Shamir Bin-Karim, Alireza Bafandeh, Christopher Vermillion

University of North Carolina at Charlotte, Charlotte, NC, USA

ARTICLE INFO

Keywords:
Bayesian Optimization
Optimal control
Energy systems
Wind energy
Airborne wind energy systems

ABSTRACT

This paper presents a framework by which a data-driven optimization technique known as Bayesian Optimization can be 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. Determining the optimal operating altitude of an AWE system is challenging, as the wind speed constantly varies with both time and altitude. Furthermore, without expensive auxiliary equipment, the wind speed is only measurable at the AWE system's operating altitude. In this work, Gaussian Process modeling and Bayesian Optimization are used in real-time to optimize the AWE system's operating altitude efficiently, without the use of auxiliary wind profiling equipment. Specifically, the underlying objective function is modeled by a Gaussian Process (GP); then, Bayesian Optimization utilizes the predictive uncertainty information from the GP to determine the best subsequent operating altitude. In the AWE application, context-dependent Bayesian Optimization is used to handle the time-varying nature of the wind shear profile (wind speed vs. altitude). Using real wind data, our method is validated against three baseline approaches. Our simulation results show that the Bayesian Optimization method is successful in significantly increasing power production over these baselines.

1. Introduction

Wind energy is one of the most promising renewable resources for displacing traditional fossil fuel sources. 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 (a kite, rigid wing, or aerostat) 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.

The vast energy resource from high-altitude winds has attracted the attention of numerous research and commercial ventures over the past two decades (Altaeros Energies Website, 2017; Ampyx Power Website, 2017; Kitenergy Website, 2017; Makani Power, 2017; SkySails GmbH Website, 2017; Windlift, 2017). 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. To-date, the majority of research in the area of AWE system control has focused on the latter problem, i.e., crosswind motion control (Canale, Fagiano, & Milanese, 2010; Fagiano, Milanese, & Piga, 2012; Isaacs, Hoagg, Hussein, & Olinger, 2011; Vermillion, Grunnagle, & Kolmanovsky, 2012; Williams, Lansdorp, & Ockels, 2007). Comparatively fewer studies have focused on the impact of adjusting altitude (Bafandeh & Vermillion, 2016; Vermillion, 2013). However, the first two papers make the assumption of a monotonic wind profile that conforms to a power law model, an assumption that greatly simplifies the altitude optimization problem but is seldom satisfied in the *instantaneous* wind shear (wind speed vs. altitude) profile (as we shall see in this paper). The more recent publication, Bafandeh and Vermillion (2016), does take into account the possibility of non-monotonic wind shear profiles but uses an extremum seeking (ES) formulation that only achieves convergence to *local*, rather than global, optima.

Only recently has any effort been undertaken to take into account the *stochastic* nature of the wind shear profile in the context of a *global* altitude optimization. An initial attempt to address this problem using model predictive control (MPC) is detailed in Bin-Karim, Bafandeh, and Vermillion (2016) and its expanded journal version (Bin-Karim, Bafandeh, Baheri, & Vermillion, submitted for publication). In the

☆ This work was supported by the National Science Foundation (NSF), Award No. 1437296.
* Corresponding author.
E-mail addresses: akhayatb@uncc.edu (A. Baheri), sbinkari@uncc.edu (S. Bin-Karim), abafande@uncc.edu (A. Bafandeh), cvermill@uncc.edu (C. Vermillion).

https://doi.org/10.1016/j.conengprac.2017.09.007
Received 5 March 2017; Received in revised form 8 September 2017; Accepted 11 September 2017
0967-0661/

Fig. 1. Altaeros Buoyant Airborne Turbine (BAT), Image Credit: Altaeros Energies Website (2017).

context of the altitude optimization problem, MPC attempts to balance exploration with exploitation in optimizing the altitude setpoint. However, this MPC formulation requires an offline characterization of the statistical properties (conditional mean and conditional variance) of the wind shear (wind speed vs. altitude) profile, thereby necessitating the collection of a very substantial amount of data offline in order to kick-start the optimization.

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 well-studied problems in the machine learning community is the design of optimization algorithms for a real-world applications using scarce data. In existing literature, this problem has been studied in the context of sequential decision-making problems where the goal is to learn the behavior of an objective function (called exploration) while simultaneously trying to maximize or minimize the objective function (called exploitation). As an efficient and systematic approach to balance exploration and exploitation, Bayesian Optimization has been applied to various real-world problems (Abdelrahman, Berkenkamp, Poland, & Krause, 2016; Baheri, Deese, & Vermillion, 2017; Baheri, Ramaprabhu, & Vermillion, 2017; Calandra, Gopalan, Seyfarth, Peters, & Deisenroth, 2014; Garnett, Osborne, & Roberts, 2010). In general, Bayesian Optimization aims to blend exploration and exploitation in such a way that it finds the global optimum of an unknown, expensive-to-evaluate, and black-box function within only a few evaluations. One popular approach is to model the unknown function as a Gaussian Process (GP), where Bayesian Optimization puts prior belief the overall structure of that objective function. At every step of Bayesian Optimization, the next operating point is selected to maximize some *acquisition function*, which characterizes (i) how much will be learned by visiting a candidate point (exploration) and (ii) what the likely performance level will be at that next candidate point (exploitation) (Brochu, Cora, & De Freitas, 2010; Snoek, Larochelle, & Adams, 2012).

In the case of the altitude optimization problem at hand, the decision variable is the operating altitude at the next time step, denoted by z_{k+1}, where k is the current time step. The objective is net power output, taking under consideration the power required to control the system at a given altitude and the power required to adjust the altitude. In this work, we focus specifically on the Altaeros Energies Buoyant Airborne Turbine (BAT), pictured in Fig. 1. The BAT maximizes power output by *hunting* for the optimal altitude.

This paper represents a greatly expanded version of our original conference publication, Baheri and Vermillion (2017), which provides an initial overview of the Bayesian Optimization approach and some initial results for the AWE application. The present paper provides additional details with regard to the Bayesian Optimization formulation and lower-level control structure that were omitted from Baheri and Vermillion (2017) for the sake of brevity. Additionally, this paper expands the results of Baheri and Vermillion (2017) through the following contributions:

- We compare the performance of Context-Dependent Bayesian Optimization (CDBO) with the three different acquisition functions introduced in Section 3.1.2.
- We present an analysis of the proposed algorithm under varying rate limits (which dictates the maximum allowable altitude change from one time step to the next), using a full week of real wind shear data.

The remainder of this paper is structured as follows: In Section 2, we summarize the wind shear profile data and formalize the objective function for the altitude optimization problem. Section 3 introduces the fundamentals of Gaussian Process (GP) modeling, Bayesian Optimization, and the proposed algorithm. Section 4 provides detailed results, indicating significant improvement in net energy generation over baseline approaches.

2. Wind shear profile, altitude adjustment, and energy generation models

2.1. Wind shear profile

To maximize generated energy, it is of interest to operate the AWE system at altitudes where the wind velocity, V_{wind}, is as close as possible to the rated wind speed, V_{rated}. Consistent with intuition, if $V_{\text{wind}} < V_{\text{rated}}$, the net power production will diminish because the turbine is producing less power. Less obviously, the net power production will also diminish whenever $V_{\text{wind}} > V_{\text{rated}}$ because more power will be expended controlling the system. While it is possible to install wind profiling equipment to continuously monitor V_{wind} as a function of altitude (z), such profiling equipment is very costly. Table 1 shows an estimated cost comparison between different wind profiling technologies, including fixed LiDAR and SODAR units, as well as a weather balloon that continually is raised and lowered through the domain of allowable altitudes. These costs account for capital expenses, along with operational expenses that include the cost of the energy required to operate the system. The cost estimate for the weather balloon system includes not only the balloon, but also the on-board instrumentation and winch system that continually actuates the weather balloon. Given these costs, it is of economic interest to use the AWE system itself to periodically explore the wind shear environment, taking into account the spatio-temporal variability of the wind in deciding which altitude to explore next. In the case of the Altaeros BAT (used as a case study in this work), the wind speed is measured through the use of an ultrasonic anemometer, as discussed in Vermillion, Grunnagle, Lim, and Kolmanovsky (2014). Anemometer measurements (which provide apparent wind speed) are used in conjunction with GPS data to compute the true wind speed. This alleviates the need for expensive equipment but only provides a wind speed measurement at the AWE system's operating altitude.

To conduct our study, we used data obtained from a Doppler radar-based wind profiler in Cape Henlopen State Park (located in Lewes, Delaware). The data includes wind speed at 30 min intervals, at 48 altitudes up to 3000 m, over the course of one year (Archer, 2014). Due to technical constraints for AWE system flight altitude, we used