

An Uncertainty-Aware Data-Driven Predictive Controller for Hybrid Power Plants

Manavendra Desai, Himanshu Sharma, Sayak Mukherjee, Sonja Glavaski

Pacific Northwest National Laboratory
Richland, WA, 99354, USA

{manavendrabalwant.desai, himanshu.sharma, sayak.mukherjee, sonja.glavaski}@pnnl.gov

Abstract

The advancements in energy storage technologies and the increasing shift toward developing renewable energy power plants to reach decarbonization goals have led to a surge of interest in *hybrid* power plants. These plants combine wind or solar power generating capacity with energy storage. However, relying on variable renewable resources such as wind and solar can introduce uncertainty, both in plant power production and in meeting the demands of the electricity power grid. To ensure the reliable and efficient operation of the hybrid power plant in a grid-connected mode, an *optimization-based predictive control strategy* can be utilized to support the coordination of power generation, energy storage, and power demand. Given the advancements in data-driven modeling for complex engineering and scientific applications, this work utilizes a data-driven predictive control method, namely *subspace predictive control*, to coordinate hybrid power plant components and meet a desired power demand despite the presence of weather uncertainties. An uncertainty-aware data-driven predictive controller is proposed and its potential is analyzed using real-world electricity demand profiles. For the analysis, a hybrid power plant with a wind, solar, and co-located energy storage capacity of 4 MW each, is considered. The analysis shows that the predictive controller can track a real-world inspired electricity demand profile despite the presence of weather-induced uncertainties and be an intelligent forecaster for HPP performance.

Introduction

Background: The world is rapidly transitioning towards increased usage of renewable energy systems to meet energy demands and achieve decarbonization goals. Advances in renewable energy generation technologies, such as solar panels and wind turbines, along with energy storage solutions, have heightened the focus on their reliable and sustainable integration with conventional energy systems (Bolinger et al. 2023). Hybrid power plants (HPP) have emerged as a promising solution, combining two or more technologies, which may include wind turbines, photovoltaic (PV) or concentrated solar power systems, energy storage, geothermal power, hydropower, biomass, natural gas, oil, coal, or nuclear power. In this study, we focus on HPPs designed for

electricity production, specifically those leveraging renewable generation, such as wind and solar, combined with energy storage (Bolinger et al. 2023). In the U.S., as of the end of 2023, solar PV and energy storage plants have led the market, with a growing presence of wind, solar, and energy storage combinations. Additionally, there is strong interest in hybridization among developers of renewable energy systems, evidenced by an increase in interconnection requests for HPPs.

Developers are also exploring advancements in power electronics for integrating renewable assets like solar, wind, and energy storage, to mimic capabilities of traditional power plants in terms of their capacity value, dispatchability, ancillary services, and reliability. The heightened penetration of renewable energy-dominant systems also motivates researchers to reconsider revenue streams for wind-based HPPs, comparable to conventional power plants (Dykes et al. 2020). Additionally, this has inspired research in simultaneously optimizing for system design and operation (i.e control co-design) (Garcia-Sanz 2019; Herber 2017; Bird et al. 2024; Sharma et al. 2024) to ensure long term efficiency and reliable operation of the system. Furthermore, new regulations for grid operators and agencies are being assessed to mandate roles of renewable-energy generators in fortifying the power grid during contingencies (Li et al. 2023). Consequently, the design objectives for HPPs are shifting from merely reducing the levelized cost of energy to identifying income streams that maximize revenue through participation in time-varying energy pricing markets, ancillary services, and capacity markets. Further discussions on HPPs are available in (Dykes et al. 2020).

Motivation: The ability of HPPs to function comparably to traditional power-generation plants is often compromised by the uncertainty and variability inherent in wind, solar, and energy storage technologies. This presents both a challenge and an opportunity for developers of HPPs to innovate control systems that deliver economic benefits akin to traditional power plants. Addressing this uncertainty can enhance the reliability of and decision-making in HPPs. With the advancement and proliferation of sensing technology in complex physical systems, data-driven predictive controller designs are anticipated to offer the necessary flexibility for HPP deployment, surpassing that of model-based controller designs. For hybrid systems encompassing multiple assets,

such as wind, solar, and storage, capturing the dynamic operations of the plant is crucial for effective controller design.

Related work: In the age of "big data", data-driven modeling and decision-making for various engineering and scientific applications intersect significantly with the field of control theory (Annaswamy, Johansson, and Pappas 2024). Both fields mutually enhance their applicability to complex engineering challenges (Brunton and Kutz 2022). Recent advances in model-free decision-making have spurred significant developments in reinforcement learning and adaptive dynamic programming, both used in Markov decision processes (MDPs) (Sutton 2018; Powell 2007), as well as in control-theoretic state-space approaches with rigorous guarantees (Vrabie et al. 2009; Jiang and Jiang 2012; Mukherjee, Bai, and Chakraborty 2021). Among these data-driven control methodologies, behavioral system theory has shown considerable promise, as highlighted in works such as (De Persis and Tesi 2019; Van Waarde et al. 2020). This progress has led the development of predictive extensions, including methods like data-enabled predictive control (Coulson, Lygeros, and Dörfler 2019) and advancements in subspace predictive control (Fiedler and Lucia 2021). A comprehensive overview of data-driven control from a control theoretical perspective is provided by (Prag, Woolway, and Celik 2022), and for process control applications, by (Tang and Daoutidis 2022). A systematic exploration of how data-driven models and approaches can leverage cutting-edge machine learning methods to identify system dynamics crucial for control design is provided in (Brunton and Kutz 2022). Lastly, the integration of learned components into traditional control techniques, or deriving control actions directly from data, is increasingly popular due to advances in computing power and data acquisition. Reinforcement learning-based approaches have been identified as promising control solutions across various engineering applications, with detailed reviews presented by (Wang et al. 2022; Hewing et al. 2020).

Objective and main contributions: This paper proposes, and investigates the potential of, an uncertainty-aware data-driven predictive controller as an intelligent decision-maker and supervisor for a hybrid power plant. In this work, intelligence of the controller is considered and demonstrated in the context of its potential ability to a) encode data-driven HPP dynamics, b) provide forecasts of the power output of each component of the HPP, and c) anticipate and control the state of charge/discharge of the energy storage device in the event of peak demands from the electricity power grid.

Next, a simulation study that utilizes real-world electricity demand profiles is presented and the performance of the controller is analyzed. Lastly, open research inquiries to make the decisions of the controller trustworthy and explainable, are discussed.

The rest of this paper is organized as follows. Section introduces the HPP considered in this work and the design and application of the uncertainty-aware data-driven predictive controller. Results of the simulation study and related discussions are provided in Sections and , respectively. Finally, concluding remarks are made in Section .




Figure 1: The control scheme of the hybrid power plant (HPP) considered in this work. The HPP supervisory controller (HSC) attempts to track a load reference P_r by selecting optimal setpoints (superscript u) for the underlying plant components. Measured outputs (superscript y) are summed to obtain the overall HPP output P_l . C and G respectively represent controller and system dynamics, for each plant component.

Methodology

Hybrid power plant: This work considers an HPP that comprises a wind farm, solar farm, and battery as an energy storage device. Figure 1 shows the *supervisory* control scheme for HPP control used in this work. Setpoints are denoted by u , and measured outputs are denoted by y . The HPP supervisory controller (HSC) selects optimal setpoints for tracking an electricity load power reference P_r . The subscripts w , s , and b , denote the relation to the wind farm, solar farm, and battery, respectively. The wind and solar farms generate power outputs, P_w^y and P_s^y , respectively. Excess power produced (i.e., when $P_w^y + P_s^y > P_r$), is used to charge the battery ($P_b^y < 0$). In the event of a deficit (i.e., when $P_w^y + P_s^y < P_r$), the battery is discharged (i.e., $P_b^y > 0$) to support load-tracking. Each plant energy generation component (solar PV, turbines) comprises a dedicated and well-tuned proportional-integral-derivative (PID) controller, C , for tracking the respective setpoints - P_w^u , P_s^u and P_b^u . The wind farm, solar farm, and battery, for tracking the reference load, are approximately modeled as first-order systems, G . Note that P_w^y and P_s^y are limited by the wind and solar energy available, thus making weather forecasts and the associated uncertainties in available wind and solar energy critical to HPP performance.

Subspace predictive control: The data-driven predictive controller introduced in this work is based on the method of *subspace predictive control* (Favoreel, De Moor, and Gevers 1999). Subspace predictive control (SPC) is a parameter-free technique that identifies a multi-step prediction model of a system using only previously measured input-output data and uses this model to formulate an optimal control problem.

Consider the measured input and output data, U and Y , respectively, to comprise T trajectories of a multi-input and multi-output dynamic system. Each trajectory is $L = T_{ini} + N$ samples long. N is the length of the prediction horizon of the controller. T_{ini} is the number of the most recent measurements of U and Y of relevance. Let $u \in \mathbb{R}^m$ and $y \in \mathbb{R}^p$. Therefore, U and Y can be ex-

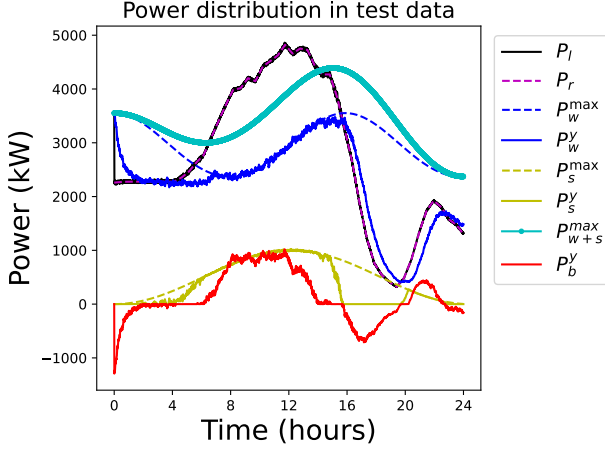


Figure 2: A sample load-tracking result. A constrained feedback-optimization based controller generates input-output data for multiple load references P_r (Ortmann et al. 2022).

pressed as $U = [u_L^1 \ u_L^2 \ \dots \ u_L^T]$ and $Y = [y_L^1 \ y_L^2 \ \dots \ y_L^T]$, where $u_L^j = [(u_L^{j,1})^\top \ (u_L^{j,2})^\top \ \dots \ (u_L^{j,m})^\top]^\top \in \mathbb{R}^{Lm}$ and $y_L^j = [(y_L^{j,1})^\top \ (y_L^{j,2})^\top \ \dots \ (y_L^{j,p})^\top]^\top \in \mathbb{R}^{Lp}$. Fulfilling Assumptions 1, 2, and 3, discussed in (Fiedler and Lucia 2021), ensures that the values of T_{ini} , N and T , satisfy observability requirements for data-driven predictive control, and make U persistently exciting. Splitting each of U and Y as $U = [U_{T_{ini}}^\top \ U_N^\top]^\top$ and $Y = [Y_{T_{ini}}^\top \ Y_N^\top]^\top$, a matrix S^* can be calculated as

$$S^* = \arg \min_S \|S \underbrace{\begin{bmatrix} Y_{T_{ini}} \\ U_{T_{ini}} \\ U_N \end{bmatrix}}_M - Y_N\|_F^2, \quad (1)$$

to give $S^* = Y_N M^\dagger$, where \dagger denotes a Moore-Penrose pseudo-inverse. Here, $U_{T_{ini}} \in \mathbb{R}^{T_{ini}m \times T}$ and $Y_{T_{ini}} \in \mathbb{R}^{T_{ini}p \times T}$. Additionally, $U_N \in \mathbb{R}^{Nm \times T}$ and $Y_N \in \mathbb{R}^{Np \times T}$.

If U and Y are sufficiently rich to cover diverse trajectories or *behaviors* of the HPP system, the matrix S^* can be used to represent any single input-output trajectory $y_{T_{ini}}$, y_N , $u_{T_{ini}}$, u_N , as

$$y_N = S^* \begin{bmatrix} y_{T_{ini}} \\ u_{T_{ini}} \\ u_N \end{bmatrix}, \quad (2)$$

where $y_{T_{ini}} \in \mathbb{R}^{T_{ini}p}$, $u_{T_{ini}} \in \mathbb{R}^{T_{ini}m}$, $u_N \in \mathbb{R}^{mN}$, and $y_N \in \mathbb{R}^{pN}$. Therefore, given the past T_{ini} sequence of input-output measurement data, (1) provides a relation between inputs and output over a future horizon of N samples. Since S^* is calculated using diverse trajectories of the HPP system, the trajectory of (u_N, y_N) is expected to be admissible by the HPP dynamics.

Data curation: In the simulation study presented in this work, the HPP is considered to serve a region with electricity demands equivalent to the combined day-long power

profiles of the Hornsea 1 wind farm, and the solar farms of PES Region 10 in East England (van der Meer 2023).

The input-output data required to determine S^* in (1) is obtained by tracking day-long demand profiles of multiple days, spread over a year. For data collection, a constrained feedback-optimization (FO) based *reactive* controller is utilized as the *HSC* of the HPP in Figure 1. The FO uses a single output measurement (i.e., y) of the most recently sampled interval to optimally select setpoints (i.e., u) for the next control action (Ortmann et al. 2022). Perfect setpoint-tracking by the inner control-loops of the HPP, and completely certain weather forecasts, are assumed. A noise-to-signal ratio of 0.02 is considered while recording y and u . Figure 2 shows a sample result for tracking an electricity load-reference P_r using the FO controller. Here, the power outputs of the wind farm (i.e., P_w^y) and solar farm (i.e., P_s^y) are upper bounded by the completely certain and maximum available wind (i.e., P_w^{\max}) and solar energy (i.e., P_s^{\max}).

Uncertainty-aware predictive controller: This work introduces an SPC-based data-driven predictive controller for supervisory control of an HPP (see figure 1), and analyzes its decision-making in the presence of uncertain weather conditions, particularly wind speed and the wind energy available to the wind farm. The controller is framed as follows.

$$\arg \min_{u_N, y_N, \sigma_u, \sigma_y} \sum_{k=1}^N \left\{ \frac{1}{2} Q_r J_T^{k,2} + \lambda \Delta P_w^{y,k,2} + \lambda \Delta P_w^{u,k,2} \right\} + \lambda_u \|\sigma_u\|^2 + \lambda_y \|\sigma_y\|^2$$

s.t. $J_T^k = P_r^k - P_w^{y,k} - P_s^{y,k}$ (3)

$$y_N = S^* \begin{bmatrix} y_{T_{ini}} \\ u_{T_{ini}} \\ u_N \end{bmatrix} + \begin{bmatrix} \sigma_y \\ \sigma_u \\ \mathbf{0} \end{bmatrix} \quad (4)$$

$$P_b^{u,k} = P_r^k - P_w^{y,k} - P_s^{y,k} \quad (5)$$

$$0 \leq P_w^{u,k}, P_w^{y,k} \leq q P_w^{\max} \quad (6)$$

$$0 \leq P_s^{u,k}, P_s^{y,k} \leq P_s^{\max} \quad (7)$$

$$P_b^{\min} \leq P_b^{u,k}, P_b^{y,k} \leq P_b^{\max}. \quad (8)$$

Here, u_N is a solution trajectory of setpoints sent by the *HSC* to the HPP components. y_N is the prediction of HPP outputs over the horizon of N samples. Q_r is a penalty on the sample-wise tracking cost J_T^k in (3). λ penalizes large changes in setpoint and power output of the wind farm. (4) constrains u_N and y_N to within the admissible system trajectories, as shown in (2). $\sigma_u \in \mathbb{R}^{T_{ini}m}$ and $\sigma_y \in \mathbb{R}^{T_{ini}p}$ are heavily penalized relaxations added in (4) to add robustness against measurement noise (Coulson, Lygeros, and Dörfler 2019). In (5), the setpoint for the battery is constrained to autonomously charge the excess ($P_b^{u,k} < 0$), or discharge to support a deficit ($P_b^{u,k} > 0$), in the combined output of the wind and solar farms. Constraints (6)-(8) set bounds on the admissible setpoints to, and power outputs of, the HPP components. $q P_w^{\max}$ is an upper bound for the available wind power in the presence of uncertain forecasts for wind speed. The uncertainty is assumed to have a normal distribution and $q P_w^{\max}$ is set using a suitable quantile value. Therefore, (6) is an individual probabilistic constraint and adds uncertainty-




Figure 3: Day-long forecasts for available wind and solar power to be used in the simulation study in Section . Uncertainty added to the wind-speed translates to uncertainty in available wind power ($P_{w,un}^{max}$).

awareness to the predictive controller. In this work, $N = 20$, $T_{ini} = 20$, $L = 40$, and $T = 1000$. Additionally, $p = 3$ and $m = 3$.

Simulation setup: The HPP in Figure 1 is simulated using the uncertainty-aware data-driven predictive controller as the *HSC*. The wind and solar farms, and the battery, have a rated capacity of 4 MW each. The battery is assumed to have a sufficient state-of-charge at all times and its cycling limits are not considered. The electricity demand profiles in (van der Meer 2023) are linearly scaled to the order of five megawatts to be consistent with the rated output of the HPP. The available wind power (i.e., P_w^{max}) and solar power (i.e., P_s^{max}) are functions of the wind speed and solar irradiance, respectively (see Figure 3). For this study, their profiles are assumed sinusoidal and are generated synthetically. To test the uncertainty-awareness of the predictive controller, a normally distributed uncertainty in wind speed, of mean 0 m/s and standard deviation 0.1 m/s, is applied. Therefore, $q_{P_w^{max}} = \mu_{P_w^{max}} + q_w \sigma_{P_w^{max}}$ in (6). Here, the quantile $q_w = -0.4$, and $\mu_{P_w^{max}}$ and $\sigma_{P_w^{max}}$ are mean and standard deviation of the uncertain wind power profile. Figure 2 shows the load-profile P_r used as the reference to the predictive controller.

The optimization problem in (3)-(8) is solved using the NLP solver from CasADi on a laptop with an 11th Gen Intel(R) Core(TM) i7-1185G7 3.00GHz processor and 16GB RAM. The mean solve time was 0.3 s, much lower than the sampling interval of 20 s of the predictive controller. Lastly, given this sampling interval, the predictive controller solves for the next $20N$ seconds using measurement data from the past $20 T_{ini}$ seconds.

Results

Open-loop control: The potential of the uncertainty-aware data-driven predictive controller is tested by first applying




Figure 4: Open-loop uncertainty-aware setpoint selection and tracking for the wind farm. $P_w^{y,ol}$ is the *open-loop* response of the wind farm to setpoints $P_w^{u,SPC}$ of the predictive controller, in the presence of uncertainty in available wind power $P_{w,un}^{max}$. The error of the predicted output $P_w^{y,SPC}$ with respect to $P_w^{y,ol}$, normalized by the mean of $P_w^{y,SPC}$, is 6.5%.

its chosen setpoints, $P_w^{u,SPC}$, $P_s^{u,SPC}$, and $P_b^{u,SPC}$, to the HPP of Figure 1, in an open-loop fashion, over the prediction horizon of N samples. Therefore, in open-loop control, the entire sequence of setpoints u_N , solved by the predictive controller in (3)-(8), is applied to the HPP without any feedback of the measured HPP output to the controller for a revised calculation of u_N . The HPP output obtained is the open-loop response. The simulation is conducted over $N = 20$ samples and with the initial sample at 11.88 hours. At this time of the day, Figure 2 shows maximum utilization of wind and solar energy, albeit in the absence of uncertainty. Therefore, such an initial condition can be used to test the uncertainty-awareness of the predictive controller since $P_w^{u,k}$ and $P_w^{y,k}$ are expected to approach an uncertain P_w^{max} .

Figure 4 plots the open-loop response of the wind farm. The setpoints $P_w^{u,SPC}$ remain below the uncertain wind power $P_{w,un}^{max}$, demonstrating the uncertainty-awareness of the predictive controller. The error of the predicted wind farm output $P_w^{y,SPC}$ with respect to the open-loop power output $P_w^{y,ol}$, normalized by the mean of $P_w^{y,SPC}$, is 6.5%. $P_w^{y,ol}$ follows $P_w^{u,SPC}$ reasonably well given that the inner control-loop for the wind farm is well-tuned. The deviation of $P_w^{y,SPC}$ from $P_w^{u,SPC}$ can be attributed to noise in the data collected to calculate S^* in (1), the relaxation terms in (4), and the input-output dynamics of the HPP captured in S^* .

Figure 5 shows the open-loop response of the solar farm, $P_s^{y,ol}$, and battery, $P_b^{y,ol}$, to the setpoints $P_s^{u,SPC}$ and $P_b^{u,SPC}$, respectively. The error of the predicted solar farm output $P_s^{y,SPC}$ with respect to $P_s^{y,ol}$, normalized by the mean of $P_s^{y,SPC}$, is 8.5%. For the battery, this value is 10%.

Figure 6 plots the predicted total power output of the HPP,




Figure 5: Open-loop setpoint selection and tracking for the solar farm and battery. $P_s^{y,ol}$ and $P_b^{y,ol}$ are respective *open-loop* responses of the solar farm and battery to setpoints $P_s^{u,SPC}$ and $P_b^{u,SPC}$ of the predictive controller. The error of the predicted solar farm output $P_s^{y,SPC}$ with respect to $P_s^{y,ol}$, normalized by the mean of $P_s^{y,SPC}$, is 8.5%. For the battery, this value is 10%.

$P_l^{SPC} = P_w^{y,SPC} + P_s^{y,SPC} + P_b^{y,SPC}$, over the prediction horizon of N samples. The error of P_l^{SPC} with respect to P_l^{ol} , normalized by the mean of P_r over the prediction horizon, is 6.1%. The uncertainty in the maximum available wind power, $P_{w,un}^{max}$, leads to non-smooth variation in P_l^{SPC} and the total available wind and solar energy $P_{w,un+s}^{SPC}$. The total power setpoint to the wind and solar farms in open-loop is $P_{w+s}^{u,ol} = P_w^{u,SPC} + P_s^{u,SPC}$. Its increasing deviation from $P_{w,un+s}^{max}$ is undesired since it is beneficial to utilize the available and renewable wind and solar energy before discharging the battery to support the tracking of P_r . Given the reasonably low normalized errors in the predictions of the data-driven controller for the open-loop outputs of the power plant components, the predictive controller demonstrated potential in providing its solutions as *forecasts* of HPP power generation.

Closed-loop control: Next, the predictive controller is tested in a closed-loop fashion. In closed-loop control, the setpoints selected by the predictive controller are applied to the HPP in a receding horizon fashion. The first setpoint of the solution u_N , $u_N^1 = \{P_w^{u,1}, P_s^{u,1}, P_b^{u,1}\}$ is applied and the measurement of the resulting output $y = \{P_w^y, P_s^y, P_b^y\}$ is feedback to the controller to solve the optimization again. Noise to signal ratio of zero is considered to isolate the response of the controller to uncertainty in the available wind power.

Figure 7 shows the effective closed-loop tracking of the load reference P_r by the uncertainty-aware predictive controller between 4am and 8am for the test data in Figure 2. The load reference P_r is tracked well by the HPP load output P_l^{cl} . The predictive controller anticipates P_r to ex-




Figure 6: Open-loop uncertainty-aware load tracking for the hybrid power plant in Figure 1. P_l^{ol} is the *open-loop* output of the hybrid power plant in response to setpoints of the predictive controller, in the presence of uncertainty in available wind power. The error of the predicted HPP output, $P_l^{SPC} = P_w^{y,SPC} + P_s^{y,SPC} + P_b^{y,SPC}$, with respect to P_l^{ol} , and normalized by the mean of P_r over the prediction horizon, is 6.1%. The combined setpoints to the wind and solar farm, $P_{w+s}^{u,ol}$, obey the probabilistic constraint in (6) to remain below the available wind and solar energy, $P_{w,un+s}^{max}$.

ceed the maximum available wind and solar energy $P_{w,un+s}^{max}$ and discharges the battery ($P_b^{y,SPC,cl} > 0$) to support load tracking. Maximum available solar energy is utilized ($P_s^{y,SPC,cl} \approx P_s^{max}$) and the wind farm output $P_w^{y,SPC,cl}$ is limited by the uncertain available wind power $P_{w,un}^{max}$.

Forecasting power plant output: Figure 8 provides the $N = 20$ samples long solution of the predictive controller at evenly spaced sampling intervals between 6am and 7am for the test data in Figure 2. At each sampling interval, the corresponding (uncertain) upper bound on available wind power, $P_{w,un}^{max}$, is plotted as a cyan dot. Given that the predictive controller showed potential to predict the open-loop performance of the HPP, its solutions can be explored as *forecasts* for the HPP outputs over a future horizon.

Figure 8 (a) provides forecasts of the setpoints to the wind farm, $P_w^{u,SPC,cl}$, each plotted as a solid line of different color. Figure 8 (b) provides forecasts of the power output of the wind farm, $P_w^{y,SPC,cl}$, each plotted as a solid line of different color. All forecasts obey the probabilistic constraint in (6) and remain below $P_{w,un}^{max}$. The underlying black dashed line in the respective plots represents the setpoint and true output recorded in the simulated closed-loop control of the HPP.

Figure 9 (a) provides forecasts of the setpoints to the battery, $P_b^{u,SPC,cl}$, each plotted as a solid line of different color. The underlying dashed line in black represents the setpoint utilized in closed-loop control of the HPP. The dashed cyan line represents the excess ($\Delta P_l > 0$) or deficit ($\Delta P_l < 0$) in power production of the wind and solar farms, relative to

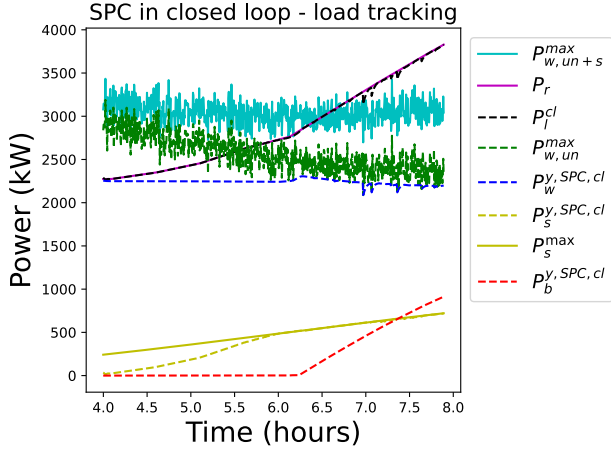


Figure 7: Uncertainty-aware closed-loop data-driven predictive control of the hybrid power plant in Figure 1. The plant output P_l tracks the load reference P_r well. The predictive controller anticipates P_r to exceed the maximum available wind and solar energy $P_{w,un+s}^{max}$ and discharges the battery ($P_b^{y,SPC,cl} > 0$) to support load tracking. Maximum available solar energy is utilized ($P_s^{y,SPC,cl} \approx P_s^{max}$) and the wind farm output $P_w^{y,SPC,cl}$ is limited by the uncertain available wind power $P_{w,un}^{max}$.

the load reference P_r . In the event of a deficit, the battery must discharge ($P_b^y > 0$) to support the load. As seen at ≈ 6.2 hours, the battery setpoints in solid orange *forecasts* the need to discharge and support the load, and turns positive toward its tail. This demonstrates the capability of the predictive controller to anticipate and control the state of the battery, despite the presence of uncertainty in the available wind power, in the event of a peak demand from the electricity power grid. Such a capability can provide data, on the future participation of the battery, to the manager of an HPP.

Figure 9 (b) provides a forecast of the total output of the HPP with respect to the load reference P_r . Each forecast is a solid line of a different color and is obtained by summing the predicted outputs, $P_w^{y,SPC}$, $P_s^{y,SPC}$, and $P_b^{y,SPC}$, over the prediction horizon. The underlying dashed line in black is the true total output obtained in the simulated closed-loop control of the HPP. Given that the true output of the power plant (in dashed black) closely follows the forecasts, the predictive controller is shown to be uncertainty-aware and capable of providing useful forecasts of power generation. This capability can give a manager of an HPP data on the anticipated behavior and output of the plant and assist in HPP management.

Discussions

The previous section showed that the data-driven predictive controller has potential for uncertainty-aware closed loop control of a hybrid power plant. Additionally, in the form of a supervisor, i.e., the *HSC* in Figure 1, the controller is shown to be capable of intelligent decision-making for a)

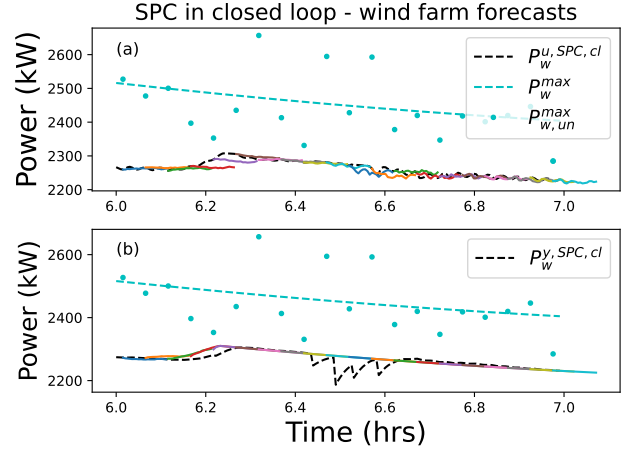


Figure 8: The solution of the uncertainty-aware data-driven predictive controller can be treated as a *forecast* or expected trajectory of the power plant output over a future horizon. The uncertain and available wind power $P_{w,un}^{max}$ is dotted at specific sample times. At these sample times, solid lines represent the prediction of future setpoints, in (a), and wind farm output, in (b). The underlying dashed lines in black represent the setpoints and true output obtained in simulated closed-loop control of the hybrid power plant.

optimal setpoint selection for each HPP component, and b) anticipation and control of the state of battery charge/discharge, despite the presence of uncertainty. Next, we discuss a few research inquiries to retain this predictive controller intelligence, and making the decisions and forecasts trustworthy and explainable.

1. How good is my data?

The uncertainty-aware predictive controller presented in this work is completely data-driven. Its ability to constrain permissible future trajectories of the system in (4) relies on the quality of the HPP input-output measurement data used to define the matrix S^* in (1). Therefore, it is necessary to develop a quantitative and/or qualita-

Table 1: Impact of uncertainty, weights, and quantiles, on the relaxations σ_y and σ_u of the behavior constraint in (4). $\|\cdot\|_F$ is the Frobenius norm. $|\Delta P_w^y|$ reflects the change in solution of the predictive controller, for wind power output $P_w^{y,SPC}$, from a baseline scenario of no uncertainty and $\lambda_u = \lambda_y = 10$. The prediction horizon considered is $N=20$ samples long and begins at 11.88 hours for the load profile P_r shown in Figure 2.

Test	$\ \sigma_u\ _F$	$\ \sigma_y\ _F$	$ \Delta P_w^y $
No uncertainty	8 kW	5 kW	0 kW
Uncertainty, $\lambda_u = \lambda_y = 10$	38 kW	17 kW	115 kW
Uncertainty, $\lambda_u = \lambda_y = 10^5$	1 kW	0.8 kW	118 kW
Uncertainty, $q_w = -0.4$	38 kW	17 kW	115 kW
Uncertainty, $q_w = -1.6$	152 kW	66 kW	660 kW

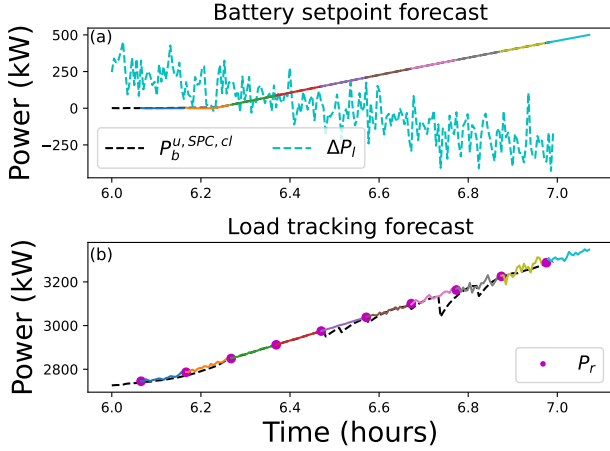


Figure 9: (a) Each solid line is a forecast of the setpoint to the battery, over the prediction horizon of the predictive controller. $\Delta P_l > 0$ is the excess ($\Delta P_l > 0$) or deficit ($\Delta P_l < 0$) in the combined power production of the wind and solar farms, relative to the load reference. In the event of a deficit, the battery must discharge ($P_b^y > 0$) to support the load. At ≈ 6.2 hours, the battery setpoints in solid orange *forecast* the need to discharge and support the load, and turn positive toward the tail of the forecast. This demonstrates the capability of the predictive controller to anticipate and control the state of the battery, despite the presence of uncertainty in the available wind power. (b) Each solid line represents a forecast of the total output of the hybrid power plant given by the predictive controller. This output is obtained by summing the predicted outputs, $P_w^{y,SPC}$, $P_s^{y,SPC}$, and $P_b^{y,SPC}$. Since the true output of the power plant (in dashed black) closely follows the forecasts, the predictive controller is shown to be uncertainty-aware and capable of providing useful forecasts of power generation.

tive metric of how well the data U and Y represent the dynamics of the HPP. If this is not the case, the most recent input-output measurements, $u_{T_{ini}}$ and $y_{T_{ini}}$, used in (4), will represent system dynamics different from that encoded by U , Y , and S^* . This can lead to

- (a) large values of relaxations σ_u and σ_y in (4) to maintain a feasible optimization, and
- (b) incorrect setpoint selection (i.e, decision making) and forecasts of HPP power output.

A data-quality metric for noise-corrupted data and data prone to privacy related obfuscation is introduced in (Mith, Morales, and Poor 2024), while a min-max data-driven optimization technique is introduced in (Huang et al. 2023) for distributional robustness against uncertain data. However, the predictive controller introduced in this work is expected to run in-house, making all required measurement data available and free of obfuscation. The question asked above seeks to determine a metric that provides analytical insight on the appropriateness of the data for predictive control and inform the need to

update the data collected, and S^* .

2. Is my data-driven optimization biased?

The potential of the data-driven predictive controller to serve as a supervisor for HPP control can lead to an evolution of the objective function and constraints utilized in (3)-(8). The design of the objective function, constraints, and the tuning of penalties such as Q_r , λ_u , and λ_y , affect the admissible system outputs, y_N , and setpoints, u_N , solved by the predictive controller over a prediction horizon. If the gains are tuned too strongly to optimize the objective function or the objective demands an uncharacteristic system trajectory, the constraint in (4) may be relaxed excessively to accommodate the solution (u_N, y_N) . This trajectory could be biased toward the objective and may not be representable by U , Y , and S^* , in (1). An example of how optimization may be optimistically biased is provided in (Hobbs and Hopenstal 1989).

3. Does constraint relaxation of the admissible system trajectories offer any clues?

Excessive relaxation of the behavior constraint in (4) can lead to increased values of σ_u and σ_y in the solution of the predictive controller. An increasing relaxation is a sign of the solution trajectory (u_N, y_N) being increasingly different than the system dynamics encoded in U , Y , and S^* , in (1). This may happen due to unaccounted uncertainty in the behavior data (old input-output measurement data of the HPP), incorrectly tuned probabilistic constraints (weather forecasts of high-variance), and improper tuning of the penalties in the objective function.

For instance, Table 1 provides the Frobenius norm of the relaxations, σ_u and σ_y , under different penalties, quantiles q_w , and uncertainty. A baseline scenario of no uncertainty (wind speed and available wind power are completely certain) is considered, with σ_u and σ_y penalized with $\lambda_y = \lambda_u = 10$. For a given set of $(\lambda_y, \lambda_u, q_w)$, $|\Delta P_w^y|$ reflects the change in solution of the predictive controller, for wind power output $P_w^{y,SPC}$, with respect to the baseline scenario. The prediction horizon considered is $N=20$ samples long and begins at 11.88 hours for the load profile P_r shown in Figure 2.

Overall, the presence of uncertainty in the weather forecasts (i.e., wind speed in this study) leads to an increased relaxation of the order of a 100 kW. This can be attributed to the following. The wind farm output is bounded by the probabilistic constraint in (6) to remain below the quantile q_w of the uncertain and maximum available wind power, as shown in Figure 7. The resulting conservative but uncertainty-aware production of wind power differs in behavior from the uncertainty-free input-output measurement data, U and Y , collected using feedback-optimization in Figure 2. Thus, the presence of a probabilistic constraint in (6), to provide uncertainty-awareness, necessitates an increased relaxation of the constraint in (4). Applying high penalties and setting $\lambda_y = \lambda_u = 10^5$, lowers relaxation of the behavior constraint. However, this comes at a cost of increased change in solution for the outputs of the solar farm and battery (not reported here for sake of brevity). Lastly, using a more

conservative quantile value of $q_w = -1.6$ to set a lower bound on the available wind power in (6), leads to a large change in the trajectory of the wind power output, leading to excessive relaxation of the constraint in (4).

Therefore, developing a threshold for admissible relaxations can provide insight into a change in underlying dynamics of the HPP. Additionally, such a threshold can point toward a need for improved tuning of penalties in the objective function, and the quantiles for the probabilistic constraints.

A starting point to determine such a threshold could be determining the relaxation required in the behavior constraint in (4) to represent the most recent HPP trajectory, $(u_{T_{ini}}, y_{T_{ini}})$. The sequence of measurements $(u_{T_{ini}}, y_{T_{ini}})$ will encode the up-to-date input-output dynamics of the HPP. Therefore, larger than expected relaxations, to describe $(u_{T_{ini}}, y_{T_{ini}})$, may point toward the need to examine the optimization framework (3)-(8) for lack of awareness of uncertainty in the data and probabilistic constraints.

Conclusion

This work proposed, and investigated the potential of, an uncertainty-aware data-driven predictive controller as an intelligent decision-maker and supervisor for a hybrid power plant (HPP). Intelligence of the controller is considered and demonstrated in the context of its ability to a) encode dynamics of an HPP through the mere use of available input-output measurement data, b) provide forecasts of the power output of each component of the HPP, and c) anticipate and control the state of charge/discharge of the energy storage device in the event of peak demands from the electricity power grid. The potential of the predictive controller was demonstrated using a simulation study that utilized real-world electricity demand profiles. The controller provided good closed-loop control, and forecasts of setpoints to and outputs of the HPP. Lastly, a few research inquiries to retain intelligence, and make the decisions and forecasts of the predictive controller trustworthy and explainable, are discussed. The next steps involve testing and analyzing the predictive controller's performance using real-world weather forecasts, higher-fidelity plant models, and joint probabilistic constraints to enhance uncertainty-awareness and evaluate its impact on solutions.

The broader objective of this work is to explore and investigate the utility of the predictive controller as an assistant to the management of an HPP, not only in producing electricity, but also in its participating in electricity markets.

Acknowledgments

The work was authored by the Pacific Northwest National Laboratory (PNNL) is operated for the DOE by Battelle Memorial Institute under contract DE-AC06-76RL01830. Funding is provided by the U.S. Department of Energy Wind Energy Technologies Office for project: Path to Nationwide Deployment of Fully Coupled Wind-Based Hybrid Energy Systems. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government.

The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

References

- Annaswamy, A. M.; Johansson, K. H.; and Pappas, G. 2024. Control for societal-scale challenges: Road map 2030. *IEEE Control Systems Magazine*, 44(3): 30–32.
- Bird, T. J.; Siefert, J. A.; Pangborn, H. C.; and Jain, N. 2024. A set-based approach for robust control co-design. In *2024 American Control Conference (ACC)*, 2564–2571. IEEE.
- Bolinger, M.; Gorman, W.; Rand, J.; and Jeong, S. 2023. Hybrid power plants: Status of operating and proposed plants. *LBNL-Report*.
- Brunton, S. L.; and Kutz, J. N. 2022. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press.
- Coulson, J.; Lygeros, J.; and Dörfler, F. 2019. Data-enabled predictive control: In the shallows of the DeePC. In *2019 18th European Control Conference (ECC)*, 307–312. IEEE.
- De Persis, C.; and Tesi, P. 2019. Formulas for data-driven control: Stabilization, optimality, and robustness. *IEEE Transactions on Automatic Control*, 65(3): 909–924.
- Dykes, K.; King, J.; DiOrto, N.; King, R.; Gevorgian, V.; Corbus, D.; Blair, N.; Anderson, K.; Stark, G.; Turchi, C.; et al. 2020. Opportunities for research and development of hybrid power plants. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States).
- Favoreel, W.; De Moor, B.; and Gevers, M. 1999. SPC: Subspace predictive control. *IFAC Proceedings Volumes*, 32(2): 4004–4009.
- Fiedler, F.; and Lucia, S. 2021. On the relationship between data-enabled predictive control and subspace predictive control. In *2021 European Control Conference (ECC)*, 222–229. IEEE.
- Garcia-Sanz, M. 2019. Control Co-Design: an engineering game changer. *Advanced Control for Applications: Engineering and Industrial Systems*, 1(1): e18.
- Herber, D. R. 2017. *Advances in combined architecture, plant, and control design*. Ph.D. thesis, University of Illinois at Urbana-Champaign.
- Hewing, L.; Wabersich, K. P.; Menner, M.; and Zeilinger, M. N. 2020. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1): 269–296.
- Hobbs, B. F.; and Hepenstal, A. 1989. Is optimization optimistically biased? *Water Resources Research*, 25(2): 152–160.
- Huang, L.; Zhen, J.; Lygeros, J.; and Dörfler, F. 2023. Robust data-enabled predictive control: Tractable formulations and performance guarantees. *IEEE Transactions on Automatic Control*, 68(5): 3163–3170.

Jiang, Y.; and Jiang, Z.-P. 2012. Computational adaptive optimal control for continuous-time linear systems with completely unknown dynamics. *Automatica*, 48(10): 2699–2704.

Li, L.; Zhu, D.; Zou, X.; Hu, J.; Kang, Y.; and Guerrero, J. M. 2023. Review of frequency regulation requirements for wind power plants in international grid codes. *Renewable and Sustainable Energy Reviews*, 187: 113731.

Mieth, R.; Morales, J. M.; and Poor, H. V. 2024. Data valuation from data-driven optimization. *IEEE Transactions on Control of Network Systems*.

Mukherjee, S.; Bai, H.; and Chakraborty, A. 2021. Reduced-dimensional reinforcement learning control using singular perturbation approximations. *Automatica*, 126: 109451.

Ortmann, L.; Maeght, J.; Panciatici, P.; Dörfler, F.; and Bolognani, S. 2022. Online Feedback Optimization for Subtransmission Grid Control. *arXiv preprint arXiv:2212.07795*.

Powell, W. B. 2007. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons.

Prag, K.; Woolway, M.; and Celik, T. 2022. Toward data-driven optimal control: A systematic review of the landscape. *IEEE Access*, 10: 32190–32212.

Sharma, H.; Wang, W.; Huang, B.; Ramachandran, T.; and Adetola, V. 2024. Multi-Objective Control Co-design Using Graph Based Optimization for Offshore Wind Farm Grid Integration. In *2024 American Control Conference (ACC)*, 2580–2585.

Sutton, R. S. 2018. Reinforcement learning: An introduction. *A Bradford Book*.

Tang, W.; and Daoutidis, P. 2022. Data-driven control: Overview and perspectives. In *2022 American Control Conference (ACC)*, 1048–1064. IEEE.

van der Meer, J. B. S. H. H. K. E. S. R. B. Y. W. D. 2023. Hybrid Energy Forecasting and Trading Competition.

Van Waarde, H. J.; Eising, J.; Trentelman, H. L.; and Camlibel, M. K. 2020. Data informativity: A new perspective on data-driven analysis and control. *IEEE Transactions on Automatic Control*, 65(11): 4753–4768.

Vrabie, D.; Pastravanu, O.; Abu-Khalaf, M.; and Lewis, F. L. 2009. Adaptive optimal control for continuous-time linear systems based on policy iteration. *Automatica*, 45(2): 477–484.

Wang, X.; Wang, S.; Liang, X.; Zhao, D.; Huang, J.; Xu, X.; Dai, B.; and Miao, Q. 2022. Deep reinforcement learning: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 35(4): 5064–5078.