Learning to Optimize: Gaussian Process Surrogates for Battery Energy Storage Coordination

Daniel Glover,*† Anup Parajuli,*† Ninad Gaikwad,*† Aryan Ritwajeet Jha,*† and Anamika Dubey*‡
* School of Electrical Engineering & Computer Science
Washington State University, Pullman, WA 99164-2752
† IEEE Graduate Student Member, ‡ IEEE Senior Member
{daniel.glover1, anup.parajuli, ninad.gaikwad, aryan.jha, anamika.dubey}@wsu.edu

***Abstract*—Utility operators face numerous challenges managing distributed energy resources (DERs) in active distribution networks due to growing uncertainty amid complex planning strategies. Providing grid services from storage assets requires solving a computationally difficult nonlinear multi time-period optimal power flow (MPOPF) using expensive solvers that suffer from modeling and measurement inaccuracies, motivating the use of surrogate modeling techniques. These approaches, however, often provide limited utility interpretability and require extensive data for development. Therefore, this work proposes a coregionalized multi-output Gaussian Process (CMOGP) surrogate to learn a fast approximation of the MPOPF with storage. The framework utilizes an inducing point strategy to indirectly capture a wide range of operational conditions with minimal training data, while simultaneously incorporating critical underlying output correlation information. Results carried out on the IEEE 123-bus test system show that the CMOGP produces near-optimal confidence-based 24-hour MPOPF solution trajectories with an average 99% computational speed-up improvement compared to a validated optimization solver in realistic operational scenarios.**

I. Introduction

A. Challenges Managing BESS in Distribution Networks

The aggressive proliferation of distributed energy resources (DERs) in active distribution networks (ADNs) presents unique operational and planning challenges for utilities. Battery Energy Storage Systems (BESS) have emerged as critical grid-edge assets, providing flexibility for distribution system operators (DSOs) to provide a variety of grid services i.e. real and reactive power support, voltage regulation, curtailment mitigation, power quality, and energy arbitrage [1], [2]. However, effective BESS management requires solving multi time-period optimal power flow (MPOPF) problems for future planning to determine setpoint dispatch strategies that balance economic and reliability objectives. These formulations are computationally complex due to their temporal coupling, large-scale nonlinear constraints, and the need to incorporate uncertainty in load and renewable forecasts [3]. In addition, the introduction of time-coupled constraints via state-of-charge (SoC) BESS dynamics increases the OPF problem dimensionality, further motivating the need for MPOPF solvers capable of predicting future system behavior under uncertainty while satisfying system and device-level constraints [4].

Traditional centralized OPF for optimizing BESS in ADNs is difficult due to the nonlinear and nonconvex nature of AC power flow equations, influencing branch flow relaxation and linearized techniques [3], [5]. Decentralized optimization schemes have been proposed to counter single point failure and reduce problem size. But these methods still face scalability concerns, require some degree of centralized oversight, and rely on a network communication infrastructure [4], [6].

B. Learning-based Optimal Power Flow (OPF) Surrogates

The computational challenges of solving distribution OPF problems has motivated the use of learning-based approaches to develop *surrogates* for optimizing ADNs [7], [8]. Deep neural networks (DNNs) have been proposed to approximate a probabilistic optimal power flow (POPF), reducing problem dimensionality [9]. However, DNNs are data hungry, *black-box* approaches, inspiring the use of statistical models such as Gaussian Processes (GP) [10]. GP surrogates have mainly been used for modeling system dynamics for designing robust control policies [11]. Tan et al. [12] focuses on standard GP surrogate models to bridge heterogeneous information exchange across difficult domains with uncertainty-aware feasibility, and [13] uses a sparse hybrid surrogate GP for ACOPF computational speedup, noting closed-form advantages. However, these works do not examine the benefits of GP surrogates for solving MPOPF problems, while appropriately addressing the convergence challenges, offering limited discussion on GP surrogates that can learn correlations in MPOPF solutions.

In this paper, we develop and analyze a coregionalized multi-output Gaussian Process (CMOGP) surrogate to learn near-optimal MPOPF solutions for centralized BESS dispatch. We train the CMOGP on validated optimal model predictive control (MPC) trajectories using a sparse inducing point strategy to discover indirect system dynamic properties while using minimal data requirements to cover a larger range of unpredictable operational scenarios. The results demonstrate that GP surrogates offer significant computational gains, enabling utilities to obtain uncertainty-aware decisions for improved BESS-enabled system operations while avoiding reliance on high-cost, high-fidelity solvers for BESS management in ADNs.

II. Model Predictive Control For BESS Optimization

A. Centralized Multi-period Optimal Power Flow

A multi-period optimal power flow (MPOPF) formulation for BESS dispatch is outlined in this section [14]. We model a

The authors greatly acknowledge support from National Science Foundation (NSF) Cyber-Physical Systems program under NSF Award Number: 2208783

979-8-3315-8138-1

single-phase radial distribution system network represented by a directed graph $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$, where $\mathcal{N}$ and $\mathcal{E}$ denote the set of all nodes and edges of the system. For each branch $(i, j) \in \mathcal{E}$, P_t^{ij} and Q_t^{ij} denote the active and reactive power flow, ℓ_t^{ij} is the squared branch current magnitude, and v_t^j is the squared nodal voltage magnitude at node j. Nodal active and reactive load demand are given by $p_{L,t}^j$ and $q_{L,t}^j$ (where $j \in \mathcal{L} \subseteq \mathcal{N}$, $\mathcal{L}$ is a set of all nodes with loads), while $p_{D,t}^j$ represents the active power injection from PV at node j (where $j \in \mathcal{D} \subseteq \mathcal{N}$, $\mathcal{D}$ is a set of all nodes with DERs), B_t^j is the battery SoC at node j (where $j \in \mathcal{B} \subseteq \mathcal{N}$, $\mathcal{B}$ is a set of all nodes with BESS), and the branch parameters r^{ij} and x^{ij} denote the resistance and reactance of line $(i, j) \in \mathcal{E}$. The MPC formulation of MPOPF problem is given as follows, where t and T are the current timestep and horizon length, respectively;

$$\min_{\{P_{c,\tau}^j, P_{d,\tau}^j\}_{\tau=t}^{t+T}} \sum_{\tau=t}^{t+T} \Bigg[(C_{sub,\tau} P_{sub,\tau} \Delta t) + \alpha \sum_{j \in \mathcal{B}} \left\{ (1 - \eta_c^j) P_{c,\tau}^j + \left(\frac{1}{\eta_d^j} - 1 \right) P_{d,\tau}^j \right\} \Bigg] \quad (1)$$

$$\text{s.t.} \quad P_\tau^{ij} = \sum_{(j,k) \in \mathcal{E}} P_\tau^{jk} + p_{L,\tau}^j - p_{D,\tau}^j - P_{d,\tau}^j + P_{c,\tau}^j \quad (2)$$

$$Q_\tau^{ij} = \sum_{(j,k) \in \mathcal{E}} Q_\tau^{jk} + q_{L,\tau}^j \quad (3)$$

$$(P_\tau^{ij})^2 + (Q_\tau^{ij})^2 = l_\tau^{ij} v_\tau^i \quad (4)$$

$$v_\tau^j = v_\tau^i - 2\left(r^{ij} P_\tau^{ij} + x^{ij} Q_\tau^{ij}\right) + \left\{(r^{ij})^2 + (x^{ij})^2\right\} l^{ij} \quad (5)$$

$$(V_{\min})^2 \le v_\tau^j \le (V_{\max})^2 \quad (6)$$

$$P_{sub,\tau} \ge 0 \quad (7)$$

$$B_\tau^j = B_{\tau-1}^j + \Delta t \eta_c^j P_{c,\tau}^j - \Delta t \left(\frac{1}{\eta_d^j} \right) P_{d,\tau}^j \quad (8)$$

$$P_{c,\tau}^j \le P_B^j, \quad P_{d,\tau}^j \le P_B^j \quad (9)$$

$$soc_{\min} B_B^j \le B_\tau^j \le soc_{\max} B_B^j \quad (10)$$

$$B_0^j = B_T^j \quad (11)$$

The MPOPF problem minimizes the energy-based cost objective in (1) conditioned on the power system operational and device-level constraints (2) - (11). This corresponds to total measured substation power P_{sub} and its cost profile C_{sub}, while minimizing BESS power loss with charging and discharging efficiency $\eta_c^j = \eta_d^j = 0.95$, and $\alpha > 0$ eliminates simultaneous charging and discharging of the battery [14]. Constraints (2) and (3) enforce active and reactive power balance at node j, and the apparent power flow rating of the line is satisfied by (4). Constraint (5) represents the branch KVL equation with voltage magnitude operational limits $(V_{\min}, V_{\max}) = (0.95, 1.05)$ p.u. in (6), and (7) prevents backflow of real power into the substation. The BESS energy dynamics are captured by (8), and charging $P_{c,t}^j$ and discharging $P_{d,t}^j$ powers are limited by the converter rating P_B^j in (9). The permissible state-of-charge (SoC) range is enforced through (10) by energy nameplate capacity B_B^j and allowable SoC limits $(soc_{\min}, soc_{\max}) = (0.3, 0.95)$. Finally, (11) requires the BESS to end with the same initial energy it begins with every $T = 24$ hours.

Table I
MPOPF MPC VALIDATION AGAINST OPENDSS

Parameters	All-time Maximum Discrepancy
Active Substation Power	0.128979 KW
Reactive Substation Power	1.039869 kVAr
Line Loss	0.030051 kW
Voltage	6.4e-05 pu

In the context of the MPOPF problem, we define $\mathbf{x}_\tau, \mathbf{u}_\tau, \mathbf{w}_\tau \; \forall j$ in the sets of $\mathcal{B}, \mathcal{L}$, and $\mathcal{D}$ as:

$$\mathbf{x}_\tau^{MPC} := \left\{ B_{\tau=t}^j \right\}, \quad (12)$$

$$\mathbf{u}_\tau := \left\{ \{P_c^j\}_{\tau=t}^{t+T}, \{P_d^j\}_{\tau=t}^{t+T} \right\}, \quad (13)$$

$$\mathbf{w}_\tau := \left\{ \{p_L^j\}_{\tau=t}^{t+T}, \{q_L^j\}_{\tau=t}^{t+T}, \{p_D^j\}_{\tau=t}^{t+T}, \{C_{sub}\}_{\tau=t}^{t+T} \right\}. \quad (14)$$

The state vector $\mathbf{x}_\tau^{MPC}$ includes only the current battery SoC B_t^j at time t (other models accept the state/observation vector $\mathbf{x}_\tau := \left[\mathbf{x}_\tau^{MPC}, \mathbf{w}_\tau\right]^T$). The disturbance vector, $\mathbf{w}_\tau$, consists of the predicted forecast profiles for system load and PV output power $\{p_L^j, q_L^j, p_D^j\}_{t=1}^T$, and the time-varying substation power cost profile $\{C_{sub}\}_{t=1}^T$ for $T = 24$ hours. The control vector $\mathbf{u}_\tau$ consists of the optimal battery charging P_c^j and discharging P_d^j powers. Once the MPOPF is solved, only the first-step control action $u_{\text{net},t}^j$ for $\tau = t$ is implemented (15), and the problem is re-solved at the next time step $t + 1$ with updated measurements and forecasts, realizing a receding-horizon control policy.

$$\{u_{\text{net}}^j\}_{\tau=t}^{t+T} = \{P_d^j\}_{\tau=t}^{t+T} - \{P_c^j\}_{\tau=t}^{t+T}, \quad \forall j \in \mathcal{B} \quad (15)$$

To validate feasibility of the MPC-based control decisions, the net BESS injections $\{u_{\text{net}}^j\}_{\tau=t}^{t+T}$ obtained from the MPOPF formulation are simulated in OpenDSS [15], comparing power flow (PF) solutions to the DSS nonlinear PF solver. Table I quantifies the maximum observed discrepancies in substation power, line losses, and voltage magnitudes, all of which are negligible (MPC has almost no plant-model mismatch), thereby confirming the physical realizability of the MPOPF-based MPC control strategy. For details see [14].

III. GP SURROGATE MODELING FRAMEWORK

A. Multi-output Gaussian Process Regression

In this section, we describe the proposed surrogate modeling framework for learning the MPOPF dispatch for BESS. Gaussian Processes (GPs) are a Bayesian class of nonparametric regression method(s) that assume a prior probability distribution over the possible function space $f_i(\mathbf{x})$ conditioned on a data set $\mathcal{D} = \{X, \mathbf{y}\}$ of observed states $X = \{\mathbf{x}_i | i = 1...n\}$ and target outputs $\mathbf{y} = \{y_i | i = 1...n\}$ [16]. The GP for a function $f(\mathbf{x})$ of random variables in the feasible space is defined by a *mean* $m(\mathbf{x}) = \mathbb{E}[f(x)]$ and *covariance* function, or *kernel*, $k(\mathbf{x}, \mathbf{x}')$, where $f(x)$ and $f(x')$ are jointly Gaussian, given as $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}\big[(f(x) - m(x)), (f(x') - m(x'))\big]$. We utilize

979-8-3315-8138-1

the *squared exponential* (RBF) kernel in (16) to capture input nonlinearities, optimizing the hyperparameters l *lengthscale* and signal *variance* σ_s^2 by taking the marginal log likelihood given by $\Theta^*(\sigma_s^2, l) = arg\max_\Theta \log pr(\mathbf{y}|X, \Theta)$.

$$k(\mathbf{x}, \mathbf{x}') = \sigma_s^2 \exp\left(\frac{-\|\mathbf{x} - \mathbf{x}'\|^2}{2l}\right) \quad (16)$$

When X is distorted with i.i.d noise, $\epsilon \sim \mathcal{N}(\mu = 0, \sigma_\epsilon^2)$, the GP assumes a Gaussian distribution over $\mathcal{D}$, and the posterior distribution conditioned on a new test point $\mathbf{x}^*$ is $\hat{f}(\mathbf{x}^*)|X, \mathbf{y}, \mathbf{x}^*, \Theta^* \sim \mathcal{N}(\hat{f}(\mathbf{x}^*), cov(f^*(\mathbf{x}^*)))$, where *cov* describes the covariance of the test point $\mathbf{x}^*$ and the prior conditioned observations $X = \{\mathbf{x}_i | i = 1...n\}$.

The multi-output (MOGP) formulation extends the standard GP to approximate $\mathbf{f} = \{f_1, ..., f_N\}^T$ output processes simultaneously considering cross-correlations, defined by a zero mean prior and multi-output covariance $k_{MO}(\mathbf{x}, \mathbf{x}') \in \mathbb{R}^{NxN}$.

$$\mathbf{f}(\mathbf{x}) \sim \mathcal{GP}(\mathbf{0}, k_{MO}(\mathbf{x}, \mathbf{x}')) \quad (17)$$

Given $\mathcal{D} = \{X, \mathbf{y}\}$ for a MOGP, the posterior distribution of $\hat{\mathbf{f}}(\mathbf{x}^*) = \{f_1(\mathbf{x}^*), ..., f_N(\mathbf{x}^*)\}$ at a new test vector $\mathbf{x}^*$ can be analytically derived across multiple outputs with a predicted mean $\hat{\mathbf{f}}(\mathbf{x}^*)$ and variance $\mathbf{\Sigma}^*$ with noise $\Sigma_I = \Sigma \otimes \mathbb{I} \in \mathbb{R}^{NxN}$

$$\hat{\mathbf{f}}(\mathbf{x}^*) = K_{MO}^{*T} \left[K(X, X') + \Sigma_I\right]^{-1} \mathbf{y} \quad (18)$$

$$\mathbf{\Sigma}^* = k_{MO}^*(\mathbf{x}^*, \mathbf{x}^*) - K_{MO}^{*T} \left[K(X, X) + \Sigma_I\right]^{-1} K_{MO}^* \quad (19)$$

where, $K_{MO}^{*T} = K_{MO}(X, \mathbf{x}^*) \in \mathbb{R}^{NNxNN}$ is a positive semi-definite matrix containing kernel vectors representing the joint covariances between X and $\mathbf{x}^*$.

B. Coregionalized MOGP Surrogate Modeling

We adopt the method of *intrinsic coregionalization* (ICM) [17] to better model the correlations between optimal MPOPF output dimensions $\mathbf{y} = \{\mathbf{u}_\tau\}_{t=1}^T$. This requires a MOGP that is *symmetric* (outputs treated equally) and *separable* ($\{X, \mathbf{y}\}$ decoupled), given $\mathcal{D}$ is *isotopic* i.e. all N outputs $\mathbf{y}$ share the same input space X in $\mathcal{D}$=$\{X, \mathbf{y}\}$=$\{\mathbf{x}_\tau, \mathbf{u}_\tau\}$. This allows for better transfer of information across regression outputs, leading to more efficient learning of correlations between different power system parameters simultaneously. We consider a separable kernel class in the form of $\mathbf{K}_{MO}(\mathbf{x}, \mathbf{x}')_{d,d'} = k(\mathbf{x}, \mathbf{x}')\mathbf{B}$ to model $D = T = 24$ output dimensions, where $\mathbf{B} \in \mathbb{R}^{DxD}$ is the coregionalization matrix encoding weighted dependencies $\mathbf{W}$ among each output, and the true covariance function is obtained as the tensor product between the RBF base covariance function $k(\mathbf{x}, \mathbf{x}')$ in (16) and $\mathbf{B}$.

$$\mathbf{B} \otimes \mathbf{K} = \begin{bmatrix} \mathbf{B}_{11} \times \mathbf{K}(\mathbf{x}, \mathbf{x}')_{11} & \cdots & \mathbf{B}_{1D} \times \mathbf{K}(\mathbf{x}, \mathbf{x}')_{1D} \\ \vdots & \ddots & \vdots \\ \mathbf{B}_{D1} \times \mathbf{K}(\mathbf{x}, \mathbf{x}')_{D1} & \cdots & \mathbf{B}_{DD} \times \mathbf{K}(\mathbf{x}, \mathbf{x}')_{DD} \end{bmatrix}$$

To validate the kernel's positive semi-definitiveness, $\mathbf{B}$ takes the form of

$$\mathbf{B} = \mathbf{W}\mathbf{W}^T + diag(\kappa) \quad (20)$$

where $\mathbf{W}$ is the weight matrix of size $[|\mathbf{y}|, |\mathbf{y}|]$ measuring output correlations, and κ is the $[|\mathbf{y}|, 1]$ diagonal vector encoding variance and designating output independence.

Representing the scalar kernel(s) $k(\mathbf{x}, \mathbf{x}')$ as $\{k_q\}_{q=1}^Q$ and a set of potential coregionalization matrices as $\{\mathbf{B}_q\}_{q=1}^Q$ for $Q = u_q^i(\mathbf{x})$ latent functions, we rewrite the model as

$$\mathbf{K}_{MO}(\mathbf{X}, \mathbf{X}) = \sum_{q=1}^{Q} \mathbf{B}_q \otimes k_q(\mathbf{X}, \mathbf{X}) \quad (21)$$

where the functions $u_q^i(\mathbf{x})$ for $q = 1, ..., Q$ and $i = 1, ..., Q$ have a zero mean and covariance $k_q(\mathbf{x}, \mathbf{x}')$ for $i = i'$ and $q = q'$ indicating functions $u_q^i(\mathbf{x})$ share the same covariance, yet remain an independent output dimension from the modeling perspective, thus forming a linear combination of random functions. With ICM, we advantageously set $Q = 1$ to reduce computational complexity of optimizing Θ^* for each dimension in combination with learning the $\mathbf{B}$ matrices.

To improve performance, scalability, and lower computational requirements of the CMOGP, we consider a finite subset of M inducing training points $\mathbf{Z} = \{\mathbf{z}_m\}_{m=1}^M \in \mathcal{D}$. These points define a set of inducing variables $\mathbf{u}_{ind}$ over $\mathcal{D}$, generated by sampling from a uniform distribution of BESS SoC states $B_t^j \sim \text{Unif}(0, 1)$ with noisy hourly forecast profiles $\{p_L^j, q_L^j, p_D^j\}_{t=1}^T$ to simulate chaotic, heterogenous input state vectors to which the MPC must respond with an optimal solution. The inducing variables $\mathbf{u}_{ind}$ define the values of the Q latent functions at the inducing inputs M as $\mathbf{u}_{ind} = Q(\mathbf{Z}) \in \mathbb{R}^{MxM}$, and the joint-distribution over the inducing variables is also Gaussian $\mathbf{y}(\mathbf{u}_{ind}) = \mathcal{N}(\mathbf{u}_{ind}|0, \mathbf{K}_{mm})$, where $\mathbf{K}_{mm}$ is the $MM \times MM$ covariance matrix evaluated at all pairs of inducing points, expressed in the ICM case by

$$\mathbf{K}_{mm}(\mathbf{X}, \mathbf{X}) = \sum_{q=1}^{Q} \mathbf{B}_q \otimes \mathbf{K}(\mathbf{Z}, \mathbf{Z}) \quad (22)$$

The objective is to factorize the posterior distribution of the function values conditioned on the inducing points, leading to a sparse, efficient CMOGP model that can learn critical correlations between the MPOPF optimal solutions with minimal observed dispatch trajectories. This strategy reduces the size of the base RBF kernel to mxm and captures a uniform distribution of operational conditions, leading to a robust and computationally inexpensive surrogate, defined in (23).

$$\mathbf{f}(\mathbf{x}_{opt}) \sim \mathcal{GP}(\mathbf{0}, \mathbf{K}_{mm}(\mathbf{X}, \mathbf{X}')) \quad (23)$$

IV. DSS Simulation & Evaluation Case Study

A. Distribution System Simulation Environment

To perform data acquisition and comparatively evaluate the CMOGP model in a realistic distribution feeder, we utilize OpenDSS to construct the modified 123-bus test system containing multiple DERs. The system includes 17 single-phase solar PV systems rated at $2.8 - 5.6$ kVA and a single large 70 kW-280 kWh rated BESS at node 50 to offset the largest 70 kW system load. A realistic commercial loadshape profile with variable noise multiplier is applied to all loads to replicate fluctuating load dynamics, which feeds a load-based cost profile representing demand cost from the substation. Realistic solar irradiance and power output profiles from [18] are applied to each PV system. We consider a realistic *quasi-static* time

979-8-3315-8138-1

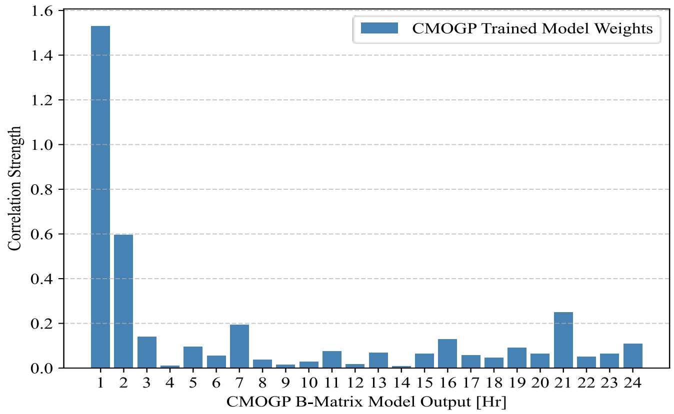

Figure 1. CMOGP 24-Hour MPOPF Solution Weight Correlation

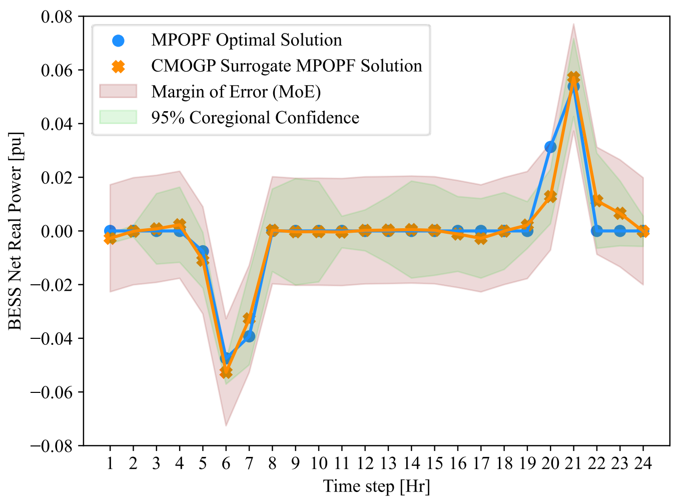

Figure 2. 24-Hour MPOPF Prediction with Coregional Confidence

series (QSTS) operational case study using an hourly time step t, where the DSO obtains system measurements with forecast predictions X and solves the 24-hour MPOPF using the validated Pyomo MPC to dispatch the updated optimal real power setpoint $\{u^j_{\text{net}}\}^{t+T}_{\tau=t}$ for $\tau = t$ to BESS. The system states update, and the simulation evolves to the next time step, repeating for N steps until the horizon is reached. All simulations are computed on a machine with Intel core i9-10900X CPU @ 3.70 GHz and 32 GB RAM.

B. Surrogate Model Evaluation & Discussion

To train the CMOGP, we set $M = 1000$ over a yearly dataset $\mathcal{D}$ to generate an inducing set of MPOPF trajectories $\mathbf{Z}$ containing state vectors X with the corresponding optimal solutions $\mathbf{y} = \{u^j_{\text{net}}\}^T_{t=1}$ generated by the MPC for $T = 24$ hours. The trained surrogate model optimizes the base RBF hyperparameters $\Theta^*(\sigma^2_s, l)$ over $\mathbf{Z}$ and learns the output matrix $\mathbf{B}$ from (20). Figure 1 shows the learned optimal solution weights $\mathbf{W}^T_{t=1}$, notably exhibiting a higher correlation strength on the first (next-hour) output 1.53, followed by the second 0.59, with less significance on the remaining 22 solution dimensions. This data provides unique insight into the model's learning of the MPOPF solution accuracy based *primarily* on the first control action to be dispatched (similar to an MPC), considering its performance impact on the other solutions.

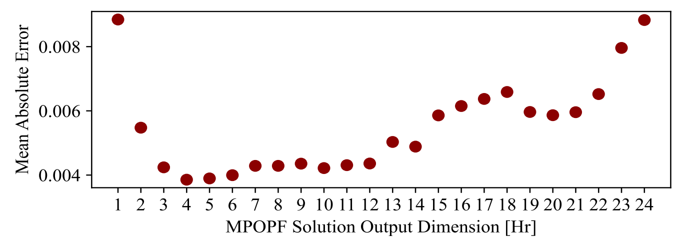

Figure 3. CMOGP MPOPF Solution Approximation Error: x1000 Cases

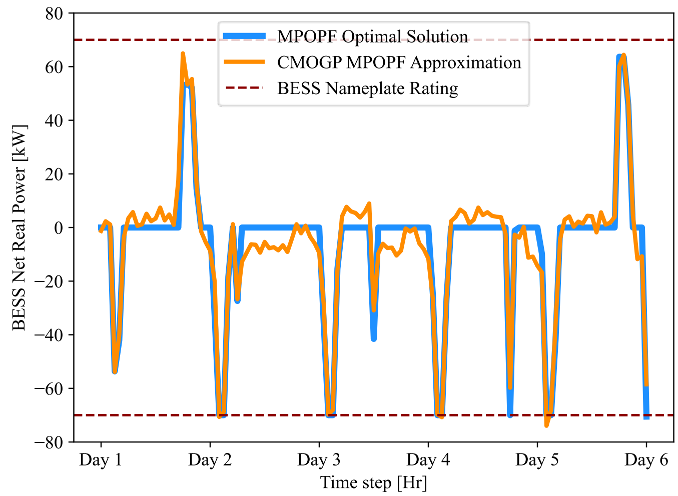

Figure 4. 5-Day Consecutive Hourly MPOPF Simulation

To evaluate the CMOGP, we first sample 1000 unseen test trajectories from $\mathcal{D}$. Figure 2 shows the complete 24-hour surrogate MPOPF prediction (orange) given a single test vector $\mathbf{x}^*$ compared with the optimal solution (blue). A main benefit of the CMOGP naturally yields a prediction-based confidence across each output dimension using the *margin of error* (MoE) $MoE = z_{val}\frac{\sigma^2_{std}}{\sqrt{M}}$, where $z_{val} = 1.96$ is the 95% confidence interval, and σ^2_{std} is the population std deviation. The MoE (red) provides a posterior region of solution uncertainty based on the entire test set distribution, but without considering output correlation. Due to the model kernel decoupling, the CMOGP advantageously provides a more accurate MPOPF confidence metric using the κ vector from $\mathbf{B}$ with the base kernel covariance, defined as the *coregional confidence* (CC), computed as $CC = \sigma^2_s \otimes \kappa$. The CC (green) provides a definitive measurement of uncertainty quantification w.r.t. each output solution dimension inside the MoE, where a tighter bound around the solution point (1 & 2) signifies a smaller variance with higher proportional weighted output confidence, compared to a wider bound. Figure 3 shows the computed mean absolute error (MAE) for all 1000 test vectors by solution dimension, indicating a slightly larger error predicting the first and last solutions of the MPOPF.

Finally, we evaluate the CMOGP in TWO realistic closed-loop QSTS 5-day hourly simulations, comparing performance of the CMOGP (orange) with the validated MPC (blue). In the first simulation, the updated $\mathbf{x}^*$ is given to the solver and CMOGP, allowing each to produce a solution to the MPOPF for comparison, and the first optimal action(s) are applied to

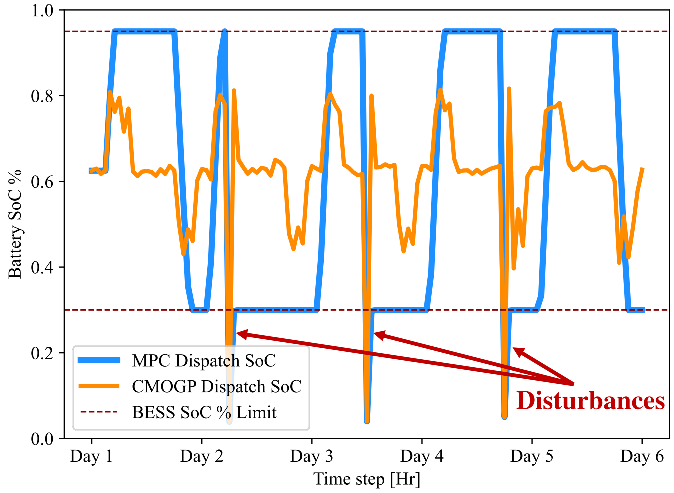

Figure 5. 5-Day Independent Soc % Hourly MPOPF Simulation

Table II
IEEE-123-BUS SYSTEM: FIVE-DAY QSTS TESTING PERIOD

MPOPF Optimal Solution Error			
Solver	**MAE**	**Min Error**	**Max Error**
CMOGP	**0.00484**	0.00011	0.0176

MPOPF Computational Speed [sec]			
Solver	**Avg Solve Time**	**Min Solve Time**	**Max Solve Time**
MPC	6.958	3.973	**231.986**
CMOGP	**0.0244**	0.016	0.044

the BESS, moving the simulation forward one hour (see Figure 4). In the second simulation, the CMOGP acts independently of the MPC, receiving a different $\mathbf{x}^*$, dispatching the next-hour near-optimal action, and moving forward to let the BESS evolve according to its previous state to monitor the SoC (see Figure 5. At hours, 30, 60, and 90, the BESS SoC is purposely disturbed (lowered) to demonstrate recoverability, constraint satisfaction, and potential computational advantages.

Figure 4 compares the optimal solutions by both models, again proving the CMOGP can produce near-optimal solutions in an evolving operational scenario. Table II shows the approximate solution MAE for the CMOGP is 0.00484 with a maximum error of 0.0176. In Figure 5, the surrogate manages the battery SoC in an acceptable manner, only briefly violating the lower bound of (10) during the disturbance, exactly as the MPC does, and immediately recovering. In fact, the CMOGP better handles constraint (11), regularly returning the SoC to its initial SoC every 24 hours. Further, the SoC disturbances create two problematic cases of *infeasibility* and *max iterations* for the MPC, causing a spike in solver runtime, as seen in Table II, unlike the surrogate, which incurs little computational cost. While the Pyomo solver maintains an average MPOFP solve time of 6.958 sec, the CMOGP approximates the MPOPF at an average time of 0.0244 sec with minimal error, a 99.6% computational improvement.

V. CONCLUSION

In conclusion, this work has demonstrated the computational and performance-based benefits of developing a learning-based GP surrogate model to provide ADN operators with a fast, reliable MPOPF approximator for BESS asset management. When compared to a validated MPC, the CMOGP proved to dispatch near-optimal BESS setpoints in a realistic QSTS operational scenario with minimal error, while mitigating disturbances and maintaining constraint satisfaction. The model also provides confidence-based decisions for improved planning under uncertainty, while offering a significant computational speedup. Regarding future work direction, we aim to focus on grid impact using a larger DER penetration to assess scalability with more complex MPOPF scenarios, extending the surrogate framework in combination with deep reinforcement learning for more robust and safety-critical control design.

REFERENCES

[1] A. Chhetri, D. K. Saini, and M. Yadav, "Applications of bess in electrical distribution network with cascading failures study–a review," *IEEE Access*, 2024.

[2] T. Gangwar, N. P. Padhy, and P. Jena, "Energy management approach to battery energy storage in unbalanced distribution networks," *IEEE Transactions on Industry Applications*, vol. 60, no. 1, pp. 1345–1356, 2023.

[3] K. Prakash, M. Ali, M. N. I. Siddique, A. A. Chand, N. M. Kumar, D. Dong, and H. R. Pota, "A review of battery energy storage systems for ancillary services in distribution grids: Current status, challenges and future directions," *Frontiers in Energy Research*, vol. 10, p. 971704, 2022.

[4] A. Dubey, S. Paudyal *et al.*, "Distribution system optimization to manage distributed energy resources (ders) for grid services," *Foundations and Trends® in Electric Energy Systems*, vol. 6, no. 3-4, pp. 120–264, 2023.

[5] L. Gan and S. H. Low, "Convex relaxations and linear approximation for optimal power flow in multiphase radial networks," in *2014 power systems computation conference*. IEEE, 2014, pp. 1–9.

[6] C. Qiu, Y. Qian, Z. Lin, and Y. A. Shamash, "Distributed optimization for the trade-off between state-of-charge balancing and strategic battery usage in a networked bess," *Journal of Energy Storage*, vol. 105, p. 114550, 2025.

[7] S. Mohammadi, V.-H. Bui, W. Su, and B. Wang, "Surrogate modeling for solving opf: A review," *Sustainability*, vol. 16, no. 22, p. 9851, 2024.

[8] R. Gupta and Q. Zhang, "Data-driven decision-focused surrogate modeling," *AIChE Journal*, vol. 70, no. 4, p. e18338, 2024.

[9] C. Srithapon, P. Fuangfoo, P. K. Ghosh, A. Siritaratiwat, and R. Chatthaworn, "Surrogate-assisted multi-objective probabilistic optimal power flow for distribution network with photovoltaic generation and electric vehicles," *IEEE Access*, vol. 9, pp. 34 395–34 414, 2021.

[10] D. Glover, P. Pareek, D. Deka, and A. Dubey, "Power flow approximations for multiphase distribution networks using gaussian processes," *arXiv preprint arXiv:2504.21260*, 2025.

[11] J. Wang and Y. Zhang, "A tutorial on gaussian process learning-based model predictive control," *arXiv preprint arXiv:2404.03689*, 2024.

[12] B. Tan, T. Su, Y. Weng, K. Ye, P. Pareek, P. Vorobev, H. Nguyen, J. Zhao, and D. Deka, "Gaussian processes in power systems: Techniques, applications, and future works," *arXiv preprint arXiv:2505.15950*, 2025.

[13] M. Mitrovic, A. Lukashevich, P. Vorobev, V. Terzija, Y. Maximov, and D. Deka, "Fast data-driven chance constrained ac-opf using hybrid sparse gaussian processes," in *2023 IEEE Belgrade PowerTech*. IEEE, 2023, pp. 1–7.

[14] A. R. Jha, S. Paul, and A. Dubey, "Spatially distributed multi-period optimal power flow with battery energy storage systems," in *2024 56th North American Power Symposium (NAPS)*. IEEE, 2024, pp. 1–6.

[15] R. C. Dugan and T. E. McDermott, "An open source platform for collaborating on smart grid research," in *2011 IEEE power and energy society general meeting*. IEEE, 2011, pp. 1–7.

[16] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[17] M. A. Alvarez, L. Rosasco, N. D. Lawrence *et al.*, "Kernels for vector-valued functions: A review," *Foundations and Trends® in Machine Learning*, vol. 4, no. 3, pp. 195–266, 2012.

[18] M. Sengupta, Y. Xie, A. Lopez, A. Habte, G. Maclaurin, and J. Shelby, "The national solar radiation data base (nsrdb)," *Renewable and sustainable energy reviews*, vol. 89, pp. 51–60, 2018.

979-8-3315-8138-1