# Strategic Investment in Energy Markets using Bayesian Optimization

Mana Jalali[1], Sina Taheri[2], and Vassilis Kekatos[3]

[1]Invenia Labs, Cambridge, CB21AW, UK; mana.jalali@invenialabs.co.uk
[2]Ibotta, Denver, CO 80202, USA; sina.taheri@ibotta.com
[3]Bradley Dept. of ECE, Virginia Tech, Blacksburg, VA 24061, USA; kekatos@vt.edu

***Abstract*—As electricity markets become more involved, strategic investment defined as the task of sizing and siting new generation units on a power system becomes increasingly computationally taxing. Adding new units in a market influences generation schedules and prices in a manner that may not be profitable for the investor overall. Strategic investment can be formulated as a mixed-integer linear program, which may not scale gracefully in larger networks and/or an increased number of stochastic scenarios. We tackle strategic investment using Bayesian optimization (BO), a machine-learning approach for minimizing hard-to-evaluate functions. BO builds a surrogate Gaussian process (GP) model for the investment cost function using function evaluations at judiciously selected points. Minimizing the surrogate cost is faster than minimizing the actual cost. By leveraging the multiparametric programming structure of strategic investment, we further expedite the standard BO approach on two fronts. First, the investment cost and its gradient with respect to the investment decision variable are evaluated more efficiently. Second, incorporating gradient information into the GP surrogate model improves its accuracy. Extensive tests on the IEEE 118-bus system corroborate the computational advantages of the BO approach over competing alternatives.**

***Index Terms*—Bayesian optimization; Gaussian processes; locational marginal prices; multiparametric programming.**

## I. Introduction

Grid modernization and the integration of renewables necessitate carefully expanding power systems and investing in new generation units. Transmission and generation planning are computationally complex problems due to uncertainties associated with load and renewable generation. Due to the coupling of investment choices and market outcomes, investment in energy markets is a non-trivial task. Further, any surplus generation modifies prices possibly in an unprofitable fashion. Thus, an investor who already owns some units in the market may end up losing money overall. To alleviate this, the investor should optimally choose the size, bids, and location of new units by maximizing its financial gain while considering the way markets are cleared by the ISO. The mentioned task is referred to as *strategic investment (SI)*.

Power systems are affected by long- and short-term uncertainties due to generation capacities, load demands, and prices. Stochastic optimization techniques have been adopted to consider these uncertainties for transmission and generation planning and investments [1], [2]. Most stochastic solutions give rise to sample approximations using several scenarios to capture diverse conditions. The SI task is posed as a bi-level optimization problem where the outer problem determines investment options upon which the lower problem clears the market [3], [4]. A typical solution is to replace the inner problem with constraints obtained from its Karush-Kuhn-Tucker (KKT). This reformulation yields a mixed-integer program (MIP), which scales unfavorably with the network size.

Multi-cut Bender's decomposition has been used to expedite the MIP-based formulation of SI in [5]. Nonetheless, the approach adds new cuts at each iteration that can increase the computational burden. Another group of SI solutions simplifies the MIP using approaches such as the consensus version of the alternating direction method of multipliers (ADMM) [6]; or progressive hedging [7], [8]. These methods, however, lack optimality guarantees. The inner problem of bi-level SI models the clearing process. A simplified formulation of this process is the DC-optimal power flow (DC-OPF). In fact, the DC-OPF must be solved for many combinations of investments and uncertain scenarios, which is computationally challenging. Reference [9] mitigated this problem by utilizing multiparametric programming (MPP). MPP uses the KKT conditions of parametric linear or quadratic programs to find closed-form expressions for optimal primal/dual variables [10], [11], [9]. Levering MPP, reference [9] puts forth two MPP-based schemes: a grid search and a stochastic gradient descent. The former becomes impractical as the number of new investment units increases; the latter can be sensitive to initialization.

This work develops a novel approach for dealing with SI using BO [12]. The approach relies on modeling the investment cost as a Gaussian process over the decision variables, and searches the optimum by iterative function evaluations [13]. Inspired by [9], function evaluations at each iteration of the BO are further accelerated using MPP. We further improve the speed and optimal investment cost by devising a derivative-aided extension of BO upon recognizing two key observations: *i)* The gradient of the investment cost with respect to the decision variables can be readily computed via MPP; and *ii)* Incorporating gradient information along with function values can enhance the accuracy of the GP surrogate model. Within the power system domain, BO has been used before to deal with hosting capacity analysis of distribution grids [14], [15]. Nevertheless, the advantages of using the MPP problem structure to further enhance BO are unique to

this work. Numerical tests compare the proposed BO-based framework against the SGP-MPP in [9], and corroborate that it is indeed more scalable compared to existing alternatives especially when the number of investment units increases.

The contributions of this work are twofold: *c1)* Adopting a surrogate probabilistic model for the SI cost, we propose a framework that exploits BO for solving SI. *c2)* We tailor the BO-based framework to a derivative-aided setup to improve upon data efficiency, which in turn yields improvements in the minimum cost and computational time.

## II. Problem Statement

An investor participating in an electricity market intends to build new generation sites. They would like to find the location and size of the new units optimally so that their overall net profit from units they already own as well as the new units, is maximized. To educate its strategic investment decision, the investor surrogates the market with a simplified model of networked-constrained economic dispatch or DC-OPF. Given load demands $\boldsymbol{\ell}$, the ISO decides generation schedules and locational marginal prices (LMPs) upon solving

$$\begin{aligned}
\min_{\mathbf{p}_r,\mathbf{p}_i} \quad & f_r(\mathbf{p}_r)+f_i(\mathbf{p}_i) && && \text{(1a)}\\
\text{s.to} \quad & \mathbf{1}^\top(\mathbf{p}_r+\mathbf{p}_i-\boldsymbol{\ell})=0 && :\lambda_0 && \text{(1b)}\\
& -\bar{\mathbf{f}}\le \mathbf{S}(\mathbf{p}_r+\mathbf{p}_i-\boldsymbol{\ell})\le \bar{\mathbf{f}} && :\underline{\boldsymbol{\mu}},\overline{\boldsymbol{\mu}} && \text{(1c)}\\
& \mathbf{0}\le \mathbf{p}_r\le \bar{\mathbf{p}}_r && :\underline{\gamma}_r,\overline{\gamma}_r && \text{(1d)}\\
& \mathbf{0}\le \mathbf{p}_i\le \mathbf{x} && :\underline{\gamma}_i,\overline{\gamma}_i && \text{(1e)}
\end{aligned}$$

where $\mathbf{p}_i$ and $\mathbf{p}_r$ are the generation schedules for the units owned by the investor and its rivals, respectively. Functions $f_r(\mathbf{p}_r)$ and $f_i(\mathbf{p}_i)$ are convex quadratic functions modeling the generation bidding costs submitted to the ISO. Power is balanced by constraint (1b). Constraint (1c) enforces line limits $\mathbf{f}$ by mapping power injections to line flows through the power transfer distribution factor matrix $\mathbf{S}$ (PTDF). Constraints (1d)–(1e) confine schedules within available capacities $\bar{\mathbf{p}}_r$ and $\mathbf{x}$ for rival and investor's units. The entries of $\mathbf{x}$ corresponding to existing units are known, whereas the entries corresponding to new units are to be found by the SI task.

Given the optimal dual variables shown on the right of the constraints in (1), the LMPs are calculated as

$$\boldsymbol{\pi}=-\lambda_0\mathbf{1}+\mathbf{S}^\top(\underline{\boldsymbol{\mu}}-\overline{\boldsymbol{\mu}}) \tag{2}$$

and the investor would be compensated by $\boldsymbol{\pi}^\top\mathbf{p}_i$ from the ISO for generating $\mathbf{p}_i$. If $g_i(\mathbf{p}_i)$ is the true cost for generating $\mathbf{p}_i$, the net revenue of the investor is $\boldsymbol{\pi}^\top\mathbf{p}_i-g_i(\mathbf{p}_i)$.

The cost of building new units can be modeled as $\mathbf{k}^\top\mathbf{x}$ for a given vector $\mathbf{k}$ consisting of the annualized cost per MW for each candidate location of the power network. All in all, the total investment cost for the investor can be cast as

$$f(\mathbf{x}):=\mathbf{k}^\top\mathbf{x}-\frac{1}{T}\sum_{t=1}^{T}\boldsymbol{\pi}_t^\top\mathbf{p}_{i,t}-g_i(\mathbf{p}_{i,t}). \tag{3}$$

The summation stems from approximating the expected value of the revenue and is over $T$ scenarios. The term *scenario* here refers to different loading and bidding conditions for (1). The terms $(\mathbf{p}_{i,t};\boldsymbol{\pi}_t)$ denote the related optimal primal/dual solutions of (1) across scenarios indexed by $t=1,\ldots,T$.

The SI task can now be formulated as

$$\begin{aligned}
\min_{\mathbf{x}\in\mathcal{X}} \quad & f(\mathbf{x}) && \text{(4a)}\\
\text{s.to} \quad & \{\boldsymbol{\pi}_t,\mathbf{p}_{i,t}\}_{t=1}^{T} \text{ being solutions of (1).} && \text{(4b)}
\end{aligned}$$

Set $\mathcal{X}$ is a simple set such as $\mathcal{X}:=\{\mathbf{x}:\mathbf{0}\le\mathbf{x}\le\bar{\mathbf{x}},\ \mathbf{1}^\top\mathbf{x}\le\delta\}$, where $\bar{\mathbf{x}}$ are possibe capacity limits per bus and $\delta$ is the total capacity budget. As in [4], we adopt three assumptions: *a1)* the network topology is known and remains constant; *a2)* the investor knows $f_r(\mathbf{p}_r)$ and $\bar{\mathbf{p}}_r$ for rival units; and *a3)* problem (1) is feasible for all $\mathbf{x}\in\mathcal{X}$. Assumptions *a1)–a2)* could can be inferred from historical data [16], while *a3)* holds if the DC-OPF is feasible without the new units.

Solving (4) is challenging due to three reasons: *i)* Constraint (4b) is an optimization problem itself [cf. (1)], referred to as the *inner problem*. The SI task of (4) is referred to as the *outer problem*; *ii)* The products between primal/dual variables of the inner problem inside (3) render the cost to be non-convex; *ii)* Evaluating the cost in (3) requires solving the inner problem for $T$ scenarios, and the actual investment benefit is approximated more accurately for larger values of $T$. It becomes apparent that minimizing or even evaluating $f(\mathbf{x})$ is computationally taxing. Considering $f(\mathbf{x})$ as a black-box, our key idea is to build a surrogate Gaussian process (GP) model for $f(\mathbf{x})$, and utilize this model to solve (4).

## III. Gaussian Processes as Surrogate Models

A GP is a random process where any finite collection of its samples forms a Gaussian random vector [17, Ch.1]. Our knowledge of $f(\mathbf{x})$ can be thought of as a random process over $\mathbf{x}$. If we have already evaluated $f$ at some $\mathbf{x}_0$, then $f(\mathbf{x}_0)$ is known without uncertainty. For other $\mathbf{x}$'s, the values $f(\mathbf{x})$ could be unknown; yet one could argue by sake of continuity that our prediction for $f(\mathbf{x})$ should be correlated to $f(\mathbf{x}_0)$ depending on the distance of $\mathbf{x}$ to $\mathbf{x}_0$.

A GP model postulated on $f(\mathbf{x})$ can be trained using function evaluations. Suppose $N$ samples of $\{\mathbf{x}_n,f(\mathbf{x}_n)\}_{n=1}^{N}$ are available. Let us collect these samples of $f$ as entries of vector $\mathbf{f}\in\mathbb{R}^N$. If $f(\mathbf{x})$ is modeled as a GP, then $\mathbf{f}$ is a Gaussian random vector. Lacking any prior information and without loss of generality, its mean value can be set to zero, while its covariance $\boldsymbol{\Sigma}$ has entries $\boldsymbol{\Sigma}_{nm}=k(\mathbf{x}_n,\mathbf{x}_m)$ for all $n,m\in\{1,\ldots,N\}$, and $k(\mathbf{x}_n,\mathbf{x}_m)$ is a covariance function expressed in parametric form such as the Matern kernel

$$k(\mathbf{x}_n,\mathbf{x}_m)=\alpha\left(1+\frac{\sqrt{5}r}{\beta}+\frac{5r^2}{3\beta^2}\right)e^{-\sqrt{5}r/\beta}+\gamma\delta_{nm} \tag{5}$$

where $r=\|\mathbf{x}_n-\mathbf{x}_m\|_2$; $\delta_{nm}$ is the Kronecker delta; and $\{\alpha,\beta,\gamma\}>0$ are parameters. Granted $\mathbf{f}\sim\mathcal{N}(\mathbf{0},\boldsymbol{\Sigma})$, parameters can be estimated using maximum likelihood from $\mathbf{f}$.

Postulating a GP model on $f$ allows to make predictions on function values for $\mathbf{x}$'s that have not been evaluated yet: Consider some $\mathbf{x}$ that has not been evaluated yet, and augment

$\mathbf{f}$ by appending the function value $f(\mathbf{x})$. Function $f$ being a GP, the augmented vector is also Gaussian as

$$\begin{bmatrix} \mathbf{f} \\ f(\mathbf{x}) \end{bmatrix} \sim \mathcal{N}\left( \begin{bmatrix} \mathbf{0} \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{\Sigma} & \mathbf{k}(\mathbf{x}) \\ \mathbf{k}^\top(\mathbf{x}) & k(\mathbf{x},\mathbf{x}) \end{bmatrix} \right) \tag{6}$$

where the $n$-th entry of $\mathbf{k}(\mathbf{x})$ is $k(\mathbf{x}_n, \mathbf{x})$ for all $n$. Since $\mathbf{f}$ and $f(\mathbf{x})$ are jointly Gaussian, the conditional PDF of $f(\mathbf{x})$ given $\mathbf{f}$ is also Gaussian with mean and variance

$$\hat{f}(\mathbf{x}) = \mu(\mathbf{x}) = \mathbf{k}^\top(\mathbf{x})\mathbf{\Sigma}^{-1}\mathbf{f} \tag{7a}$$

$$\sigma^2(\mathbf{x}) = k(\mathbf{x},\mathbf{x}) - \mathbf{k}^\top(\mathbf{x})\mathbf{\Sigma}^{-1}\mathbf{k}(\mathbf{x}). \tag{7b}$$

The conditional mean $\mu(\mathbf{x})$ is the minimum mean squared error estimate of $f(\mathbf{x})$ given $\mathbf{f}$, while $\sigma(\mathbf{x})$ quantifies the uncertainty of this estimate. It is worth stressing that (6)–(7) hold for any $\mathbf{x}$. Therefore, function $\hat{f}(\mathbf{x})$ can be used as a surrogate in lieu of the hard-to-evaluate $f(\mathbf{x})$. Minimizing $\hat{f}(\mathbf{x})$ rather than $f(\mathbf{x})$ is considerably easier as the former enjoys an analytical form and does not require solving $T$ DC-OPFs. However, if $\hat{f}(\mathbf{x})$ is an inaccurate estimate, its minimum will not provide (near) optimum investment. The GP model for $f(\mathbf{x})$ can be iteratively refined by evaluating $f$ at more points. Sampling should be accomplished in a way to reduce uncertainty $\sigma(\mathbf{x})$, but also focus on areas with smaller values of $\mu(\mathbf{x})$ as we are looking for the minimizer of $f(\mathbf{x})$. BO discussed next proposes an approach to balance this trade-off.

## IV. Investment via Bayesian Optimization

The investor has already evaluated $f$ at $N$ points as $\{\mathbf{x}_n, f(\mathbf{x}_n)\}_{n=1}^N$. Let $f_N^*$ denote the smallest value amongst these $N$ evaluations, and $\mathbf{x}_N^*$ be the corresponding capacity. The investor can improve upon $f_N^*$ by evaluating $f(\mathbf{x})$ at a new $\mathbf{x}$ featuring a potentially smaller cost. Due to the high cost of evaluating $f(\mathbf{x})$, the investor has to meticulously select the next sampling point. A useful criterion is the improvement of the optimal cost calculated so far

$$\mathrm{I}(\mathbf{x}) := [f_N^* - f(\mathbf{x})]_+ \tag{8}$$

where $[z]_+ := \max\{z, 0\}$. Ideally, the investor is looking for generation capacities that maximize $\mathrm{I}(\mathbf{x})$. Nonetheless, maximizing $\mathrm{I}(\mathbf{x})$ is as expensive as minimizing $f(\mathbf{x})$. Fortunately, thanks to the surrogate GP model, the investor can replace $f(\mathbf{x})$ in (8) by $\hat{f}(\mathbf{x})$. Because $\hat{f}(\mathbf{x})$ is random, the investor should maximize the expected value of the improvement

$$\mathrm{EI}(\mathbf{x}) := \mathbb{E}\left[ [f_N^* - \hat{f}(\mathbf{x})]_+ \right] \tag{9}$$

where the expectation is with respect to the PDF in (7). Therefore, the next point to sample can be found as

$$\mathbf{x}_n = \arg\max_{\mathbf{x}\in\mathcal{X}} \ \mathrm{EI}(\mathbf{x}). \tag{10}$$

Maximizing $\mathrm{EI}(\mathbf{x})$ is much easier than minimizing $f(\mathbf{x})$ as after some mundane algebraic manipulations, the former can be shown to enjoy the closed-form expression [13]

$$\mathrm{EI}(\mathbf{x}) = \Delta(\mathbf{x}) \cdot \Phi\left( \frac{\Delta(\mathbf{x})}{\sigma(\mathbf{x})} \right) + \sigma(\mathbf{x}) \cdot \phi\left( \frac{\Delta(\mathbf{x})}{\sigma(\mathbf{x})} \right) \tag{11}$$

where $\Delta(\mathbf{x}) := f_N^* - \mu(\mathbf{x})$, while $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and CDF for the standard normal distribution $\mathcal{N}(0,1)$.

Having identified $\mathbf{x}_n$ from (10), the investor can evaluate $f(\mathbf{x}_n)$ and either stop or use the additional sample to update the surrogate model for $f$. An algorithm can be formed by iterating across three steps: *s1)* A new $\mathbf{x}_n$ is identified by (10); *s2)* The cost is evaluated for $\mathbf{x}_n$, yielding $f(\mathbf{x}_n)$, and the new sample is added to the training data; and *s3)* The surrogate model is updated using the new datum per (7) with $\mathbf{k}(\mathbf{x})$ and $\mathbf{f}$ augmented by one entry. At each iteration, the hyper-parameters of the covariance function are updated using maximum likelihood estimation. The iterations terminate when the function evaluation budget is met. Upon completion, the investor selects the sample $\mathbf{x}_n$ attaining the smallest $f(\mathbf{x}_n)$.

The initial points $\{\mathbf{x}_n\}_{n=1}^N$ are typically sampled uniformly. This helps to avoid local minima, which is as an advantage of Bayesian optimization. Assuming a total budget of $N+M$ function evaluations, the investor must solve $T(N+M)$ DC-OPFs, which remains a computationally daunting task. The issue can be alleviated by:*i)* speeding up the evaluations of function $f(\mathbf{x})$; and *ii)* improving data efficiency so the surrogate model $\hat{f}$ learns $f$ using fewer samples. Interestingly, both goals can be achieved using neat properties of MPP as discussed next.

## V. Multiparametric Programming (MPP) for Bayesian Optimization (BO)

If bidding costs for the investor and rivals are convex quadratic functions

$$f_i(\mathbf{p}_i) = \tfrac{1}{2}\mathbf{p}_i^\top \mathbf{H}_i \mathbf{p}_i + \mathbf{c}_i^\top \mathbf{p}_i \ \text{ and } \ f_r(\mathbf{p}_r) = \tfrac{1}{2}\mathbf{p}_r^\top \mathbf{H}_r \mathbf{p}_r + \mathbf{c}_r^\top \mathbf{p}_r$$

the DC-OPF of (1) becomes a quadratic program (QP). During SI and for each evaluation of $f(\mathbf{x})$, this QP has to be solved $T$ times for different bids, demands, and generation capacities. Problem (1) can be seen as a multiparametric QP (MPQP)

$$\min_{\mathbf{y}} \ \ \tfrac{1}{2}\mathbf{y}^\top \mathbf{H}\mathbf{y} + (\mathbf{C}\boldsymbol{\theta} + \mathbf{d})^\top \mathbf{y} \tag{12a}$$

$$\text{s.to} \ \ \mathbf{D}_{\text{in}}\mathbf{y} \le \mathbf{E}_{\text{in}}\boldsymbol{\theta} + \mathbf{b}_{\text{in}} \qquad : \boldsymbol{\lambda} \tag{12b}$$

$$\mathbf{D}_{\text{eq}}\mathbf{y} = \mathbf{E}_{\text{eq}}\boldsymbol{\theta} + \mathbf{b}_{\text{eq}} \qquad : \boldsymbol{\mu} \tag{12c}$$

over parameters $\boldsymbol{\theta} := [\mathbf{c}_r^\top \ \mathbf{c}_i^\top \ \boldsymbol{\ell}^\top \ \bar{\mathbf{p}}_r^\top \ \mathbf{x}^\top]^\top$ if the quadratic components $(\mathbf{H}_i, \mathbf{H}_r)$ of the bidding costs are assumed constant. Here $\mathbf{y}$ is the vector of primal variables, while $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ are the dual variables related to linear (in)equality constraints.

As $\boldsymbol{\theta}$ varies, the solution set of (12) features a particularly appealing structure [18]: The set of $\boldsymbol{\theta}$'s for which (12) is feasible can be partitioned into distinct so-termed *critical regions*, which are polytopes in $\boldsymbol{\theta}$. Within each region, the same inequality constraints in (12b) are satisfied with equality (binding), while optimal primal/dual variables are affine functions of $\boldsymbol{\theta}$. For example, region $k$ can be defined as $\mathcal{C}_k := \{\boldsymbol{\theta} : \mathbf{A}_k^\theta \boldsymbol{\theta} + \mathbf{b}_k^\theta \le \mathbf{0}\}$, and the solution to (12) can be expressed as $\mathbf{y} = \mathbf{A}_k^y \boldsymbol{\theta} + \mathbf{b}_k^y$, $\boldsymbol{\lambda} = \mathbf{A}_k^\lambda \boldsymbol{\theta} + \mathbf{b}_k^\lambda$, and $\boldsymbol{\mu} = \mathbf{A}_k^\mu \boldsymbol{\theta} + \mathbf{b}_k^\mu$. The affine parameters $(\mathbf{A}, \mathbf{b})$ vary per region and depend on the binding constraints.

The key benefit of the aforesaid structure is that when the investor solves (12) for a particular $\boldsymbol{\theta}$, it can readily chart an entire critical region $\mathcal{C}_k$ and compute the related affine parameters $(\mathbf{A}_k, \mathbf{b}_k)$. For any other $\boldsymbol{\theta}$ falling in $\mathcal{C}_k$, the investor does not need to solve (12) since the optimal primal/dual solutions are known analytically. When the function is to be evaluated at a new $\mathbf{x}$, the investor needs to solve $T$ instances of (12), yet they all have the same $\mathbf{x}$ as part of $\boldsymbol{\theta}$. On the contrary, when $f$ is evaluated over the initial $N$ points, the investor needs to solve $NT$ instances of (12), over which the $\mathbf{x}$ part of $\boldsymbol{\theta}$ will be varying. Either task can be expedited by the MPQP structure as the investor needs to actually solve the DC-OPF of (12) fewer than $T$ or $NT$ times.

Nonetheless, using MPP to speed up function evaluations is not the key novelty of this work; we have proposed this idea before in [9]. Where the MPP structure really shines is towards improving data efficiency while training the GP surrogate $\hat{f}$. *Can we train $\hat{f}$ to learn $f$ using fewer function evaluations?* The answer can be on the affirmative if any known property of $f$ is properly fused in $\hat{f}$. Our novel idea is to train the GP surrogate model to mimic not only $f(\mathbf{x})$, but also its gradient $\nabla_{\mathbf{x}} f(\mathbf{x})$. To do so, two technical questions are in order: *q1)* Can we compute $\nabla_{\mathbf{x}} f(\mathbf{x})$? and *q2)* How to incorporate function gradient data while training $\hat{f}$?

Question *q1)* can be handled efficiently thanks to MPP. Recall from (3) that $f(\mathbf{x})$ involves products $\mathbf{p}_{t,i}^{\top}\boldsymbol{\pi}_{t,i}$ between generation schedules and prices. From (2), prices are linear functions of the optimal dual variables $(\boldsymbol{\lambda}_t, \boldsymbol{\mu}_t)$ of the DC-OPF. Per MPP structure, the optimal primal/dual solutions of (1) and (12) are affine functions of $\boldsymbol{\theta}$, and thus, affine functions of $\mathbf{x}$. Therefore, the $\mathbf{p}_{t,i}^{\top}\boldsymbol{\pi}_{t,i}$ terms in $f(\mathbf{x})$ are quadratic functions of $\mathbf{x}$. Their gradients are affine, and can be computed at almost no additional cost. If the generation costs $g_i(\mathbf{p}_{t,i})$ are quadratic or affine, then $\nabla_{\mathbf{x}} f(\mathbf{x})$ is affine in $\mathbf{x}$.

Regarding *q2)*, GPs exhibit the interesting property that the derivative of a GP is a GP itself [17]. Therefore, if $f(\mathbf{x})$ is modeled as a GP, its gradient $\dot{\mathbf{f}}(\mathbf{x}) := \nabla_{\mathbf{x}} f(\mathbf{x})$ is also a GP. Moreover, its cross-covariance with $f$ can be computed as

$$\mathbb{E}[\dot{\mathbf{f}}(\mathbf{x}) f(\mathbf{x}')] = \nabla_{\mathbf{x}} \mathbb{E}[f(\mathbf{x}) f(\mathbf{x}')] = \nabla_{\mathbf{x}} k(\mathbf{x}, \mathbf{x}').$$

Its autocovariance can be similarly computed as

$$\mathbb{E}[\dot{\mathbf{f}}(\mathbf{x}) \dot{\mathbf{f}}(\mathbf{x}')^{\top}] = \nabla^2_{\mathbf{x}\mathbf{x}'} \mathbb{E}[f(\mathbf{x}) f(\mathbf{x}')] = \nabla^2_{\mathbf{x}\mathbf{x}'} k(\mathbf{x}, \mathbf{x}').$$

Thereby, if $k(\mathbf{x}, \mathbf{x}')$ is known, the covariances of $\dot{\mathbf{f}}(\mathbf{x})$ are readily available with no additional hyperparameter tuning. In a nutshell, we can easily augment the dataset from $\{\mathbf{x}_n, f(\mathbf{x}_n)\}_{n=1}^N$ to $\{\mathbf{x}_n, f(\mathbf{x}_n), \dot{\mathbf{f}}(\mathbf{x}_n)\}_{n=1}^N$. The same holds during the subsequent $M$ function evaluations. Samples $\{f(\mathbf{x}_n), \dot{\mathbf{f}}(\mathbf{x}_n)\}_{n=1}^N$ are jointly Gaussian with known covariances, and so, the model for $\hat{f}(\mathbf{x})$ in (7) can be augmented accordingly to incorporate gradient data. Gradient data can reveal the behavior of $f(\mathbf{x})$ in a neighborhood around $\mathbf{x}$.

## VI. Numerical Tests

The SI approach was tested on the IEEE 118-bus benchmark. All tests were performed on an AMD Core i7 @ 2.9 GHz (16 GB RAM) computer, using Python on Jupyter Notebook. The DC-OPFs were solved using the `ECOS` solver [19] on the `CVXPY` library. GP parameters were tuned using the `gpr` function from the `sklearn` library. The expected improvement was maximized using the `slsqp` optimizer of the `Scipy` library. The covariance of the surrogate model was selected as a Matérn-5/2 kernel [17, Ch. 4].

TABLE I
Runtime and SI cost by MPP-SGD for 5 initializations

| # of units: 2 | Initialization | | | | |
|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 |
| $f^*$ [\$/$H$] | $-5,274$ | $-5,278$ | $-5,240$ | $-5,263$ | $-5,264$ |
| Runtime [s] | 2,789 | 3,004 | 496 | 2,120 | 2,858 |

TABLE II
SI cost using BO and DA-BO for different $N$ and $N + M = 50$.

| # of units: 2 | $N = 6$ | $N = 10$ | $N = 17$ | $N = 34$ |
|---|---|---|---|---|
| $f^*$ with BO [\$/$H$] | $-5,281$ | $-5,284$ | $-5,276$ | $-5,273$ |
| $f^*$ with DA-BO [\$/$H$] | $-5,282$ | $-5,282$ | $-5,283$ | $-5,294$ |

We used hourly bidding and load data from the day-ahead PJM market for 2018, yielding $T = 8,760$ different loading and bidding conditions. The available 21 load profiles were randomly assigned to buses. Each load profile was perturbed by adding a uniformly distributed deviation of $\pm 5\%$ independently over time and buses. Load profiles were scaled, so their annual peak matched the benchmark values. For new generation units, we assumed a purchase and installation cost of $3 \cdot 10^6$ \$/MW, plus $1 \cdot 10^6$ \$/MW for operation and maintenance over 25 years. Converting that cost to dollars per hour per unit of active power for a base of 100 MVA yielded $k = 1,826.5$ in (3).

The investor already owned a generator at bus 1 and aimed at investing in two new units on buses $\{29, 95\}$. The BO method was compared against the SGD-MPP of [9]. Since SGD-MPP solution is affected by the initial point, the algorithm was run for 5 random initializations. For SGD-MPP, the step size and convergence tolerance were set to $10^{-4}$ and $10^{-6}$. Table I reports the runtimes and costs. Different runtimes and costs confirm that SGD-MPP may find local optima. Further, the times reported in Table I do not reflect the preprocessing time needed for choosing step size and error tolerance. To compare the cost achieved by SGD-MPP and the BO-based method, SI was solved using a total budget of $N + M = 50$ function evaluations. We tested four combinations of $N$ and $M$. The initial $N$ points were uniformly chosen by sampling every $\{3, 6, 10, 16\}$ points across the two-dimensional grid formed by the two new units. Table II confirms the advantage of DA-BO over SGD-MPP.

The runtime for BO (including generating training data) was 4 hours when $N = 6$ and $M = 44$. The total time for other $N$ was less since the surrogate model had to be retrained fewer times. Evidently, SGD-MPP is faster than BO as reported in Table I. However, a fair runtime comparison would be the

TABLE III
RUNTIME FOR ATTAINING OPTIMAL COST WITH $1\%$ ADDITIVE ERROR USING BO AND DA-BO FOR DIFFERENT $N$ WITH $N + M = 50$

| # of units: 2 | $N = 6$ | $N = 10$ | $N = 17$ | $N = 34$ |
|---|---|---|---|---|
| cost with BO $[\$/H]$ | $-5,264$ | $-5,284$ | $-5,274$ | $-5,265$ |
| time with BO [s] | $4,001$ | $5,190$ | $6,993$ | $11,766$ |
| cost with DA-BO $[\$/H]$ | $-5,263$ | $-5,262$ | $-5,282$ | $-5,267$ |
| time with DA-BO [s] | $3,972$ | $5,214$ | $6,984$ | $11,764$ |

TABLE IV
SI COST AND RUNTIME FOR $T = 50$ SCENARIOS USING BO FOR DIFFERENT $N$ WITH $N + M = 50$

| # of units: 2 | $N = 6$ | $N = 10$ | $N = 17$ | $N = 34$ |
|---|---|---|---|---|
| cost with BO $[\$/H]$ | $-2,704$ | $-5,293$ | $-4,170$ | $-4,151$ |
| time with BO [s] | 78 | 75 | 74 | 65 |

time BO attained the same cost as SGD-MPP. Table III reports the times at which the proposed schemes reached $99\%$ of the minimum cost along with the sub-optimal costs. The results in Table III show that if a budget lower than $50$ function evaluations is considered, the runtime of BO is comparable to SGD-MPP while attaining lower cost.

Finally, we compared BO with multiparametric programming with equilibrium constraints (MPEC) from [9]. The investment cost was evaluated using $T = 50$, since implementing MPEC for more scenarios is time-consuming. Using such a small $T$ has the implication that the cost is not properly approximated and hence the found capacities can be suboptimal. To make a fair comparison, the same $T = 50$ scenarios were used for function evaluations of the BO algorithm. The BO was run with $N = 34$ and $M = 16$. Here we report the average cost and runtime of the $100$ Monte Carlo tests using MPEC. Each test randomly selected $T$ scenarios for evaluating $f(\mathbf{x})$. The average time for solving the SI with MPEC for $100$ Monte Carlo tests was reported as $500$ seconds, while the average optimal cost was over $-4,350$ $[\$/H]$. Given the optimum $\mathbf{x}$, the cost was calculated using $T = 8,760$ *out-of-sample* scenarios. The costs and total runtimes for BO are reported in Table IV. We can see that except for $N = 6$, the BO yields better or similar results as MPEC in much less time. Evaluating $f(\mathbf{x})$ with a larger number of scenarios can improve the cost. The MPEC becomes impractical for larger $T$. For instance, solving the SI for the IEEE $118$-bus using $T = 100$ scenarios takes nearly four days, while BO takes only $160$ seconds.

## VII. CONCLUSIONS

This work put forth a novel approach for solving SI by leveraging a surrogate GP model for the hard-to-evaluate cost function. The Bayesian optimization toolbox has been tailored to the investment problem. Employing MPP accelerated solving the inner problem of SI. Leveraging a unique property of GPs allowed using the derivatives to further improve data efficiency for the BO-based approach, while derivatives were readily found through MPP. Comparing the proposed method with existing approaches on the IEEE $118$-bus system proved its superiority for obtaining a lower cost. This work sets the foundation for several exciting and practically relevant research directions. Exploiting the BO approach for efficiently solving contingency analysis, using active learning for SI, and finding the market equilibrium in presence of multiple strategic investors are interesting and pertinent future work.

## REFERENCES


[1] J. Alvarez Lopez, K. Ponnambalam, and V. H. Quintana, "Generation and transmission expansion under risk using stochastic programming," *IEEE Trans. Power Syst.*, vol. 22, no. 3, pp. 1369–1378, Jul. 2007.

[2] J. H. Roh, M. Shahidehpour, and L. Wu, "Market-based generation and transmission planning with uncertainties," *IEEE Trans. Power Syst.*, vol. 24, no. 3, pp. 1587–1598, Aug. 2009.

[3] A. J. Conejo, L. B. Morales, S. J. Kazempour, and A. S. Siddiqui, *Investment in Electricity Generation and Transmission: Decision Making under Uncertainty*, 1st ed. Switzerland: Springer, 2016.

[4] S. J. Kazempour, A. J. Conejo, and C. Ruiz, "Strategic generation investment using a complementarity approach," *IEEE Trans. Power Syst.*, vol. 26, no. 2, pp. 940–948, May 2011.

[5] S. Bhattacharjee, R. Sioshansi, and H. Zareipour, "Benefits of strategically sizing wind-integrated energy storage and transmission," *IEEE Trans. Power Syst.*, vol. 36, no. 2, pp. 1141–1151, Mar. 2021.

[6] V. Dvorkin, J. Kazempour, L. Baringo, and P. Pinson, "A consensus-ADMM approach for strategic generation investment in electricity markets," in *Proc. IEEE Conf. on Decision and Control*, Miami Beach, FL, USA, Dec. 2018, pp. 780–785.

[7] Y. Liu, R. Sioshansi, and A. J. Conejo, "Multistage stochastic investment planning with multiscale representation of uncertainties and decisions," *IEEE Trans. Power Syst.*, vol. 33, no. 1, pp. 781–791, Jan. 2018.

[8] F. D. Munoz and J.-P. Watson, "A scalable solution framework for stochastic transmission and generation planning problems," *Computational Management Science*, vol. 12, no. 4, pp. 491–518, Oct 2015.

[9] S. Taheri, V. Kekatos, and H. Veeramachaneni, "Strategic investment in energy markets: A multiparametric programming approach," *IEEE Trans. Power Syst.*, vol. 37, no. 4, pp. 2590–2600, Jul. 2022.

[10] F. Borrelli, A. Bemporad, and M. Morari, "Geometric algorithm for multiparametric linear programming," *Journal of Optimization Theory and Applications*, vol. 118, no. 3, pp. 515–540, Sep. 2003.

[11] P. Tondel, T. A. Johansen, and A. Bemporad, "An algorithm for multi-parametric quadratic programming and explicit MPC solutions," *Automatica*, vol. 39, no. 3, pp. 489 – 497, 2003.

[12] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, Jan. 2016.

[13] P. I. Frazier, "A tutorial on Bayesian optimization," 2018. [Online]. Available: https://arxiv.org/abs/1807.02811

[14] X. Geng, L. Tong, A. Bhattacharya, B. Mallick, and L. Xie, "Probabilistic hosting capacity analysis via Bayesian optimization," in *Proc. IEEE PES General Meeting*, Washington, DC, Jul. 2021, pp. 1–5.

[15] B. Barracosa, J. Morin, H. Dutrieux Baraffe, J. Bect, E. Vazquez, and J. Fournel, "A Bayesian approach for the optimal integration of renewable energy sources in distribution networks over multi-year horizons," in *International Conference and Exhibition on Electricity Distribution*, vol. 2021, Sep. 2021, pp. 2437–2441.

[16] C. Ruiz, A. J. Conejo, and D. J. Bertsimas, "Revealing rival marginal offer prices via inverse optimization," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 3056–3064, Aug. 2013.

[17] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006.

[18] Y. Ji, R. J. Thomas, and L. Tong, "Probabilistic forecast of real-time LMP via multiparametric programming," in *International Conference on System Sciences*, Kauai, HI, USA, Jan. 2015, pp. 2549–2556.

[19] A. Domahidi, E. Chu, and S. Boyd, "ECOS: An SOCP solver for embedded systems," in *European Control Conference (ECC)*, Zurich, Switzerland, Jul. 2013, pp. 3071–3076.