

# BAYESIAN OPTIMIZATION FOR AUTOMATIC TUNING OF MODEL PREDICTIVE CONTROLLERS

Leonardo D. González and Victor M. Zavala

Department of Chemical and Biological Engineering  
University of Wisconsin – Madison, Madison, WI 53706

## Abstract

Recent improvements in data collection and computation have led to an increased adoption in data-driven optimization methods like Bayesian optimization (BO) across many disciplines. However, there are still problems, such as model predictive control (MPC) tuning, that are mainly optimized via heuristics or trial-and-error-based approaches, both of which can be time-consuming; these methods are also not scalable and cannot be applied to high dimensional problems. Recent work has shown that replacing these techniques with a more systematic and scalable algorithm like BO can significantly reduce search times and deliver reliable performance. However, traditional BO methods tend not make use of any of the ever-increasing preexisting information available for physical systems. We propose a new Bayesian optimization framework that can incorporate this information via a reference model to further improve the performance of the algorithm both by further accelerating the search and reducing the probability of converging to a local solution. We apply our approach to the MPC tuning problem of back-off terms for energy storage tanks at a central heating, ventilation, and air conditioning (HVAC) plant. We observe that the incorporation of the reference model results in a 33% reduction in the computation time required to find a solution.

## Keywords

Bayesian optimization, MPC, HVAC

## Introduction

Model predictive control (MPC) is a versatile tool that has become widely adopted across various industrial sectors. While obtaining quality process models is key to ensuring the success of an MPC application, its performance can also be impacted by various hyperparameters such as the control horizon, constraint back-off terms, and the weights of the cost objectives, see Yamashita et al. (2016); Koller et al. (2018). Determining the optimal values for these settings, however, can be challenging as they typically affect performance in complex and often non-intuitive ways; as a result, users typically rely on heuristics or trial-and-error methods to set them. Typically, these approaches do not present any issues if the number of parameters being tuned is small or if changes in the performance of the application can be measured quickly. However, this is often not the case as measuring performance may require the use of simulations which take on the order of hours to run and these simulations can involve several hyperparameters.

Black-box optimization methods are a set of algorithms that are utilized to optimize systems that lack an explicit model that maps inputs to outputs as well as derivative in-

formation, as discussed in Conn et al. (2009). Several of these approaches, such as particle swarm optimization and genetic algorithms, have been used to solve the MPC tuning problem. Issues encountered when using these methods have included slow convergence and sensitivity to the initial guess, see Kolda et al. (2003). As a result, there has been a push to identify or develop black-box optimization algorithms for tuning MPC applications that are more sample-efficient and capable of consistently locating the global solution, see Forgione et al. (2019); Garriga and Soroush (2010). One of these algorithms, Bayesian optimization (BO), has proven to be an especially effective tool for hyperparameter tuning in the deep-learning setting, as shown by Snoek et al. (2015). The ability of BO to incorporate information on both model uncertainty and performance as well as its flexibility in accommodating mixtures of continuous and discrete inputs have made it one of the most effective black-box optimization algorithms; for a detailed discussion see Brochu et al. (2010); Shariari et al. (2016).

Because BO is specifically designed to handle problems that lack a system model and where input/output effects are not well understood, it is purely data-driven and does not make use of any preexisting information that may exist. However, many systems often have some form of preexisting information available and incorporating this into the optimization routine could potentially improve its performance.

---

<sup>1</sup> Corresponding author: Victor M. Zavala (E-mail: zavalatejeda@wisc.edu).

The incorporation of various types preexisting information has already been shown to significantly improve the performance of machine learning models being used for a variety of tasks such as design of experiments (Häse et al. (2020)), learning partial differential equation solutions (Raissi et al. (2019)), control (Eugene et al. (2020)), and bioengineering (Zhang et al. (2020)). In the context of BO, we believe that the best way to capture this information is in the form of a reference model that provides an approximation of the general system trends. Access to this model would allow the algorithm to identify areas of interest immediately, resulting in a more targeted search that requires less iterations to locate a solution. Additionally, one of the criticisms of data-driven methods is that they are usually not capable of handling features like constraints or feasibility considerations. A reference model can be designed to convey these features to the algorithm, keeping it from sampling in undesired regions.

We have explored the use of BO with a reference model for tuning an MPC application, see Lu et al. (2021). We have demonstrated that an effective reference model can be generated can using various approaches and is therefore extendable to several applications. Additionally, we also provide insight over why this approach proves to be effective. We see that when the BO algorithm has a reference model, it is able to quickly focus in on regions indicated to be optimal rather, resulting in a significantly more targeted search. Additionally, we also determined that by using a reference model and shifting to learning the residual rather than objective function, we are able to decrease the complexity of the function that the algorithm has to learn, which leads to improved surrogate model estimates. We applied our method to tuning of a real-world MPC application for a central heating, ventilation, and air conditioning (HVAC) plant. The goal of the algorithm is to select a set of back-off terms,  $\beta$ , that minimize the closed-loop cost of the plant. The purpose of these terms is to minimize the occurrence of constraint violations in the hot and cold thermal energy storage (TES) tanks by reserving a fraction of the storage capacity to serve as a buffer for absorbing errors in the MPC forecast model. Setting these parameters can be challenging as an explicit function relating the back-off terms and the year-long closed-loop cost is not available, additionally setting these terms via a trial-and-error/search-based strategy would be costly as simulating this system over a year involves solving over 8,700 optimization problems and requires over two hours of wall-clock time per simulation. The results we obtained demonstrate that Bayesian optimization is an effective tool for dealing with this class of problems and that when we combine BO with a reference model we can significantly improve the performance of the algorithm. While traditional BO required approximately 30 hours to locate an optimal set of  $\beta$ 's, our approach took only 6; after taking into account the time used to generate the reference model (14 hours), this amounts to a 33% reduction in the required computational time.

## Bayesian Optimization with Reference Models

We can formulate the goal of the MPC tuning problem as:

$$\min_{\xi} f(\xi) \quad (1a)$$

$$s.t. \xi \in \Omega \quad (1b)$$

where  $f(\xi)$  is the objective or cost function and  $\xi \in \Omega$ , are the MPC application hyperparameters;  $\Omega \subset \mathbb{R}^d$  is the design space across which we search for a solution. Typically,  $f$  is of the form  $f(\xi) = \sum_{i=1}^{N_f} w_i f_i(\xi)$  where  $f_i$  represents the  $i^{th}$  objective and  $w_i$  is a user-specified weight that measures its relative importance. Generally, finding a solution to (1) is challenging because an explicit relationship between the function inputs and outputs is not available. As a result, the cost function is treated as a black-box with every  $\xi$  of interest requiring an evaluation to determine its corresponding value. Additionally, sampling from the objective can be significantly expensive, limiting the number of samples that can be taken to locate a solution. Therefore, any algorithm proposed for solving this problem must be capable of exploring the design space in a systematic and efficient manner that enables it to find an optimal set of hyperparameters using a minimal number of samples.

Bayesian optimization is an algorithm designed for optimizing complex black-box functions. Given a set of  $n$  input/output observations  $\mathcal{D} = \{\xi_{1:n}, f(\xi_{1:n})\}$ , BO trains a surrogate model that provides an estimate of  $f$ . The most frequently used surrogate model, the Gaussian process (GP), assumes that the output data have a prior multivariate normal distribution  $f(\xi_{1:n}) \sim \mathcal{N}(\mathbf{m}(\xi), \mathbf{K})$  where  $\mathbf{m}(\xi) \in \mathbb{R}^d$  is the mean function and  $\mathbf{K} \in \mathbb{R}^{d \times d}$  is the covariance matrix. A positive-definite kernel function,  $k(\xi, \xi')$ , is used to build the covariance matrix such that  $\mathbf{K}_{ij} = k(\xi_i, \xi_j)$ . There is a wide array of options for the kernel function and selection is largely based on which choice of  $k(\xi, \xi')$  fits the data best; for our work we have selected to use the Matérn kernel with a smoothness parameter  $\nu = 5/2$ . The GP estimates the value of  $f$  at some new point  $\xi$  by assuming that  $f(\xi)$  is jointly Gaussian with the observed output data,  $f(\xi_{1:n})$ , resulting in a posterior normal distribution with moments:

$$\mu = \mathbf{K}(\xi, \xi_{1:n}) [\mathbf{K}(\xi_{1:n}, \xi_{1:n}) + \sigma_n^2 \mathbf{I}]^{-1} \xi_{1:n} \quad (2a)$$

$$\Sigma = \mathbf{K}(\xi, \xi) + \sigma_n^2 \mathbf{I} - \mathbf{K}(\xi, \xi_{1:n})^T [\mathbf{K}(\xi_{1:n}, \xi_{1:n}) + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{K}(\xi_{1:n}, \xi) \quad (2b)$$

where  $\sigma_n$  measures the noise of the observations and  $\mathbf{I}$  is the corresponding identity matrix. These parameters provide values for the estimated model performance,  $\mu$ , and the accompanying model uncertainty,  $\Sigma$ . This estimation of  $\Sigma$  is one of the features that really sets BO apart from other optimization methods. By incorporating it into a utility function that is used to select a new sample point, commonly referred to as the acquisition function (AF), the algorithm can be made to sample from regions that exhibit high model uncertainty, known as exploration, as well as high performance, known as exploitation. This feature forces the algorithm to search

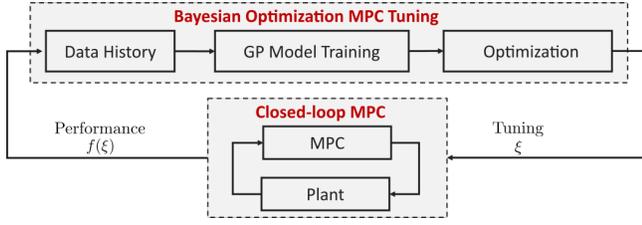


Figure 1: Block-flow diagram for implementation of BO for an MPC tuning problem

through distinct regions of  $\Omega$ , even after finding a potential solution, increasing the odds of finding the global solution. Additionally, the AF can be constructed in a manner that allows the exploration/exploitation tendencies of the algorithm to be tuned to suit the needs of the application. In this work we used the lower confidence bound (LCB) as our AF:

$$LCB(\xi) = \mu(\xi) - \kappa\sigma(\xi) \quad (3)$$

Here,  $\kappa \in \mathbb{R}_+$  is defined as the exploratory parameter and serves as a weight that determines the emphasis on exploration. Small values of  $\kappa$  will result in the BO algorithm being more focused on exploitation and large values will cause it to focus more on exploration. The exploratory parameter can either remain constant throughout the BO run or can be dynamically set so that it is large during the initial iterations to promote exploration and gradually decreases to focus on exploitation at later iterations. Regardless of the approach used, proper tuning of  $\kappa$  is essential for ensuring a balanced algorithm that does not over-explore and unnecessarily samples from suboptimal regions, or begins exploiting too quickly and misses the global solution. A new sampling point,  $\xi_{n+1}$ , is selected via minimization of the AF. Figure 1 provides an graphical representation of the process for applying BO to an MPC tuning problem.

Reference models,  $g(\xi)$ , can be used to incorporate pre-existing information into the BO algorithm essentially warm-starting the optimization routine by providing an initial approximation of the objective function. As a result, the algorithm does not have to utilize as many samples as it normally would to build up its own approximation and can quickly move to exploring regions that the reference model indicates are promising. The reference model can be built using various means such as mechanistic laws, empirical correlations, or a mix of both. In this work, we use a more coarse simulation model that takes significantly less to run than the full model while still capturing the general trends of the system. In order to incorporate  $g(\xi)$  into the BO algorithm, we first decompose our performance function as follows:

$$f(\xi) = g(\xi) + \varepsilon(\xi) \quad (4)$$

where  $\varepsilon(\xi)$  is the residual model and measures the mismatch between the reference model and the objective function. We shift the BO algorithm from learning  $f(\xi)$  to learning  $\varepsilon(\xi)$  which we approximate using a GP,  $\varepsilon(\xi) \sim \mathcal{GP}(m(\xi), k(\xi, \xi'))$ . When using a GP model, the mean function is typically set to 0, meaning that function values are drawn from  $\mathcal{N}(0, \Sigma)$ . While this can be achieved for any

function being modeled by normalizing the output data, if  $g(\xi)$  is taken to be a ground truth, it then seems more reasonable to model  $\varepsilon$  using this distribution, thereby making our method a more intuitive approach. Finally, the form of the AF in (3) must also be modified to ensure that we are optimizing over estimates of the objective function and not the residual function. Estimates for the average,  $\mu(\xi)$  and uncertainty  $\sigma(\xi)$  of  $\varepsilon(\xi)$  at any point  $\xi$  are provided by the GP model while  $g(\xi)$  is assumed to be a deterministic model and can be said to be distributed according to  $\mathcal{N}(g(\xi), 0)$ . Combining (4) with the closure of normal distributions under linear operations, for some input,  $\xi$ , we obtain

$$g(\xi) + \varepsilon(\xi) \sim \mathcal{N}(g(\xi), 0) + \mathcal{N}(\mu(\xi), \sigma(\xi)^2) \quad (5a)$$

$$\begin{aligned} & \mathcal{N}(g(\xi), 0) + \mathcal{N}(\mu(\xi), \sigma(\xi)^2) \\ &= \mathcal{N}(g(\xi) + \mu(\xi), \sigma(\xi)^2) \end{aligned} \quad (5b)$$

$$f(\xi) \sim \mathcal{N}(g(\xi) + \mu(\xi), \sigma(\xi)^2) \quad (5c)$$

This allows us to update the LCB AF to accommodate the reference model:

$$LCB(\xi) = g(\xi) + \mu(\xi) - \kappa\sigma(\xi) \quad (6)$$

### MPC Tuning for an HVAC Plant

The goal of the MPC application at a central HVAC plant is to minimize utility consumption while ensuring that heating and cooling demands are met. The plant mainly uses three utilities, water, natural gas, and electricity; the purchase costs of water and natural gas are fixed but electricity prices vary throughout the day and can be significantly higher at peak hours relative to non-peak hours. The MPC formulation and HVAC plant model can be found in Kumar et al. (2020). In order to shift production to non-peak hours, the HVAC facility is equipped with hot and cold water thermal energy storage tanks. The MPC uses a forecast model to predict demand and set production targets when electricity prices are low. However, forecasting demand is difficult resulting in frequent modeling errors which can result in TES tank constraint violations where production targets are either too high and the tanks overflow or too low and they dry up. To reduce the occurrence of these events, a set of back-off terms is introduced that reserves an upper and lower fraction of the hot and cold water tanks to provide a buffer that can absorb the impacts of forecasting errors; a graphical representation of this approach can be seen in Figure 4. Selecting appropriate values for the back-off terms is essential to ensure that they fulfill their intended purpose. If the values are too small, the buffer will not be large enough to absorb forecasting errors and constraint violations will continue to occur. Conversely, if the values are too large, they will significantly limit the amount of production that can be shifted to off-peak hours, forcing the facility to purchase more electricity during peak hours. Typically, these parameters, labelled as  $\beta_j$ , where  $\beta_j \in [0, 0.5]$  and  $j \in \{cw, hw\}$  corresponds to the cold and hot TES tanks respectively, are set manually using a combination of trial-and-error and heuristics-based approaches. However, these methods can require a significant

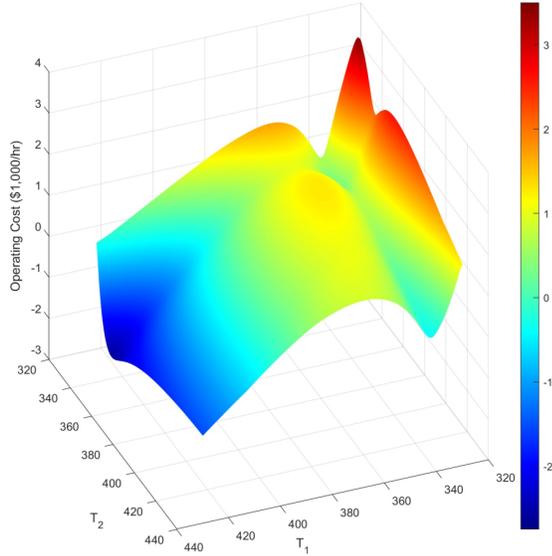
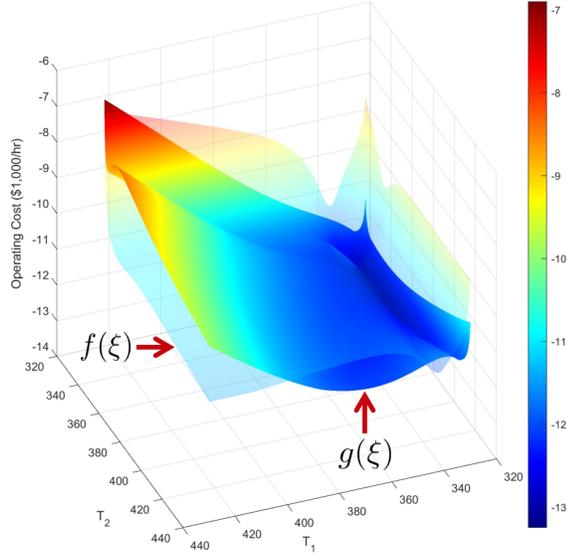


Figure 2:  $f(\xi)$  and  $g(\xi)$  for a 2-D system (top) and the corresponding residual,  $\varepsilon(\xi)$  (bottom)

number of year-long closed-loop simulations, making them computationally intensive.

We set the year-long annual closed-loop cost as our objective function,  $f(\xi)$ , and the cold and hot water back-off terms as our input parameters,  $\xi = [\beta_{cw}, \beta_{hw}]$ . Due to the fact that a model mapping back-off settings to the annual cost is not available, a simulation must be performed to determine  $f$  at every  $\xi$  of interest. Additionally, because every simulation requires approximately two hours of wall-clock time to run, we also wish to minimize the number of samples required to identify a solution. The reference model is obtained using a shortened prediction horizon simulation. In the full simulation, the prediction horizon is set to 168 hours as this appeared to be the periodicity at which electricity prices and load demand fluctuated. We determined that by reducing this horizon to 24 hours, we could decrease the simulation time

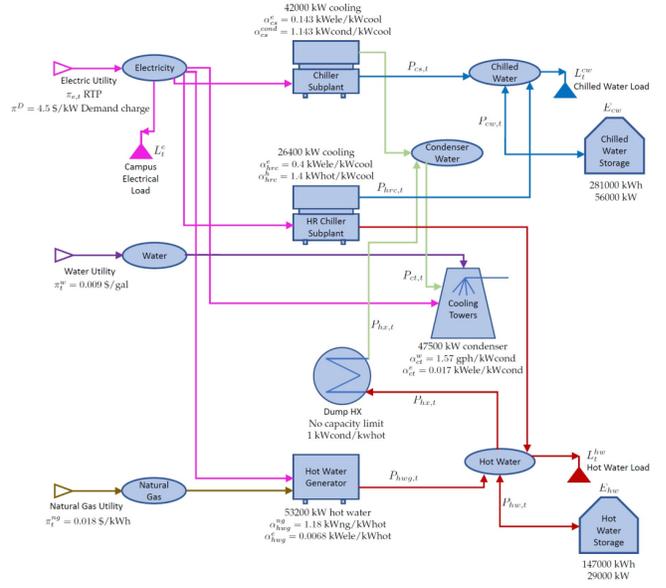


Figure 3: Schematic diagram of central HVAC plant

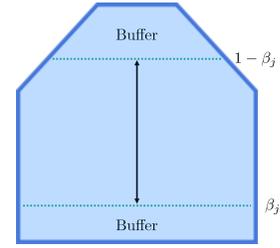


Figure 4: Cold (cw) and hot (hw) water TES tank with buffer fraction set by the back-off term,  $\beta_j \in [0, 0.5]$ ,  $j \in \{cw, hw\}$

by over 50% without significantly compromising the accuracy of the output. We use this coarser simulation model to run an instance of the BO algorithm that is heavily focused of exploration and collects 21 samples; the GP model trained at the end of the BO run is used as  $g(\xi)$ . Figure 5 provides a graphical comparison between the true objective and the reference model.

The convergence plot shown in Figure 6 illustrates the performance of traditional and reference model-guided BO. Both algorithms are able to quickly locate solutions that improve upon the baseline cost (black dashed-line) demonstrating the BO can serve as an effective tool for tuning the hyperparameters of MPC applications. When comparing the two algorithms, it is clear that our approach performs significantly better, requiring only three iterations to locate the optimal region, while traditional BO requires 15. After taking into account the time spent creating the reference model (14 hours), reference model-guided BO took about 20 hours to converge while BO alone took approximately 30, this amounts to reduction in wall-clock time of 33%. As previously mentioned, the driving force behind the performance gaps is the initial model structure provided by  $g(\xi)$ . Figure 5 the clearly demonstrates that the reference model approximates the objective well and indicates that the cost is minimal in a region near where the true global minimum exists. This causes the algorithm to move to quickly sample from this region; in fact, Figure 6 indicates that BO samples almost ex-

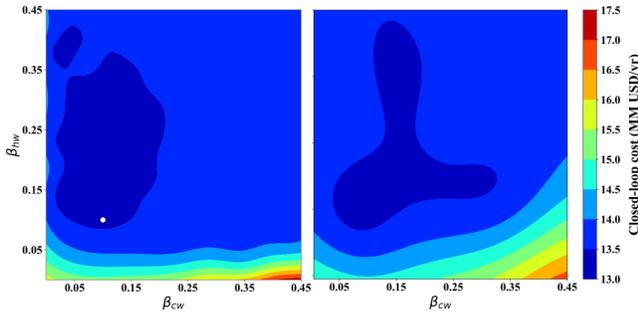


Figure 5:  $f(\xi)$  (left) and  $g(\xi)$  (right) for the MPC tuning problem. The white marker in the left plot is the location of the baseline cost for the expert-selected back-off terms as seen in Kumar et al. (2020).

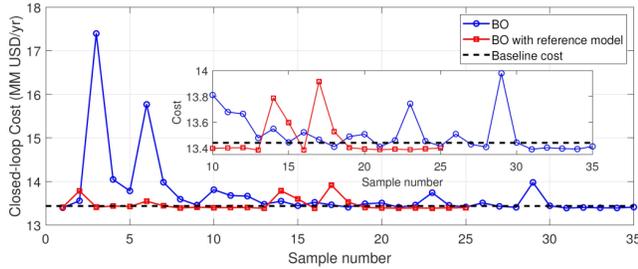


Figure 6: Iteration number vs annual closed-loop cost for traditional and reference model-guided BO with comparison to baseline cost. Close-up view provides visual confirmation that solutions obtained improve upon the base cost.

clusively from the region where  $g(\xi) \leq 13.5$  MM USD. The plots showing the distribution of samples in Figure 7 confirm this to be true: only five of the samples are taken from outside of this region. Of the remaining samples, the vast majority are taken from lower left portion of this region which is the area closest to the true solution. By comparison, the sampling distribution is significantly more dispersed when traditional BO is used; a significant number of samples are taken at or near the design space boundaries. This is because initially, the algorithm must build a more descriptive surrogate model that has enough structure and shape to point out regions that could be exploited. By preloading BO with  $g(\xi)$ , we are able to mostly bypass this initial exploratory phase and move directly onto sampling from high-potential regions.

In addition to making BO faster and more sample-efficient,

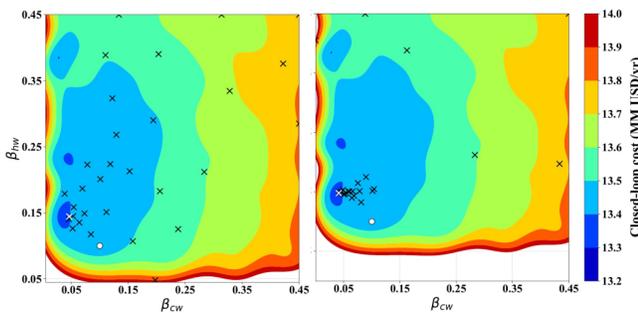


Figure 7: Distribution of function evaluations for traditional BO (left) and reference model-guided BO (right)

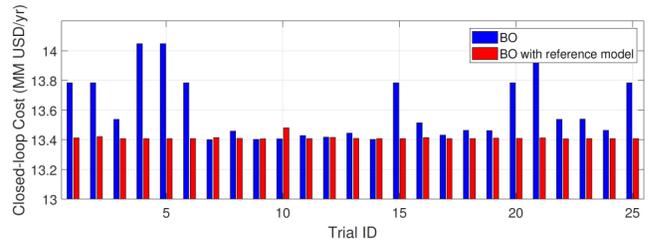


Figure 8: Optimal closed-loop annual cost for traditional and reference model-based BO across 25 different initialization points on a  $5 \times 5$  grid of the design space.

incorporation of a reference model also makes the algorithm significantly more robust. While BO was designed to be a global optimizer, like other optimization methods, it is sensitive to the initialization point. While the degree of this sensitivity can be lessened by selecting an appropriate set of hyperparameters (kernel length scale and AF exploratory parameter), tuning these can be challenging as data is required to do so. As previously mentioned, a potential solution is to make these parameters, dynamic, progressively modifying them as the algorithm progresses and data becomes available. Almost all BO/GP modeling software, including the package we used (Scikitlearn’s `gaussian_process`), are set to re-learn the kernel parameters when more data is made available. However, this still leaves the challenge of selecting the appropriate range for these hyperparameters and how many optimizer restarts to use; additionally, how and when to change the exploratory parameter is a task left largely to the user. We have observed that the use of  $g(\xi)$  significantly decreases the sensitivity of the BO algorithm to the initial guess without requiring extensive hyperparameter tuning. Figure 8 demonstrates the value of the best solution found by traditional and reference model-guided BO when initialized at different points on a 5 by 5 grid. One can clearly see that even at initialization points where traditional BO clearly fails to locate the global optimum, reference model-guided BO is still successful while using the same value for  $\kappa$  and having the same upper and lower bounds for the kernel parameters. In fact, our method delivered very consistent performance, converging essentially the same solution every run. Again, this is largely due to the fact that BO can use  $g(\xi)$  to quickly identify potentially high-value regions. In addition to providing this benefit, the reference model also captures the structure of the coarser features of  $f(\xi)$ , thereby significantly reducing the complexity of  $\epsilon(\xi)$  and simplifying the learning task of the algorithm. These results allow us to be more confident that the solution returned by our method is highly likely to be the global optimum.

## Conclusions

We have developed a framework for incorporating pre-existing system information into the BO algorithm via a reference model. This information can be in the form of mechanistic models, empirical correlations, historical data, or simplified simulation models. BO uses this model to quickly identify and sample from potentially high-value regions, re-

sulting in a significantly more targeted search. Additionally, we shift the learning task of the algorithm from learning the objective function to learning the residual function which captures the mismatch between the objective and the reference model. We observe that the structure provided by the reference model significantly reduces the complexity of the residual, making it easier to learn.

We applied this framework to an MPC tuning problem where we wished to select the set of back-off terms for a pair of thermal energy storage tanks that minimized that annual closed-loop cost at a central HVAC plant. Due to the lack of an explicit model relating these terms to the cost, the system must be simulated at every input of interest. The duration and computational intensity of the simulation makes manual or grid search approaches very expensive, forcing us to turn to more efficient algorithms like Bayesian optimization. We establish that BO is capable of solving MPC tuning problems within a reasonable time-frame, requiring only 30 hours to find a solution that improved upon an expert-selected one. Moreover, we clearly demonstrate that when we combine BO with a reference model (from a coarser and faster simulation model), we see a significant improvement in performance when compared to traditional BO: the algorithm locates the region containing the global optimum in only 20 hours (a 33% reduction in computational time). We also observe that our version of BO is significantly less sensitive to the initialization point than traditional BO, converging to essentially the same solution regardless of where it takes its first sample. Our results show that the incorporation of a reference model into BO make it a fast, sample-efficient, and robust global optimization algorithm that consistently delivers excellent results. Looking towards the future, we would like to test our framework on higher dimension problems such as MPC tuning with a larger number of tuning parameters to determine potential areas for improvement. Additionally, we would also like to investigate the minimal quality of the reference model needed to provide the algorithm with enough guidance to maintain the improvements we have observed.

## References

- Brochu, E., V. M. Cora, and N. De Freitas (2010). A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.
- Conn, A. R., K. Scheinberg, and L. N. Vicente (2009). *Introduction to derivative-free optimization*. SIAM.
- Eugene, E. A., X. Gao, and A. W. Dowling (2020). Learning and optimization with Bayesian hybrid models. In *2020 American Control Conference (ACC)*, pp. 3997–4002. IEEE.
- Forgione, M., D. Piga, and A. Bemporad (2019). Efficient calibration of embedded mpc. *arXiv preprint arXiv:1911.13021*.
- Garriga, J. L. and M. Soroush (2010). Model predictive control tuning methods: A review. *Industrial & Engineering Chemistry Research* 49(8), 3505–3515.
- Häse, F., L. M. Roch, and A. Aspuru-Guzik (2020). Gryffin: An algorithm for Bayesian optimization for categorical variables informed by physical intuition with applications to chemistry. *arXiv preprint arXiv:2003.12127*.
- Kolda, T. G., R. M. Lewis, and V. Torczon (2003). Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review* 45(3), 385–482.
- Koller, R. W., L. A. Ricardez-Sandoval, and L. T. Biegler (2018). Stochastic back-off algorithm for simultaneous design, control, and scheduling of multiproduct systems under uncertainty. *AIChE Journal* 64(7), 2379–2389.
- Kumar, R., M. J. Wenzel, M. N. ElBsat, M. J. Risbeck, K. H. Drees, and V. M. Zavala (2020). Stochastic model predictive control for central HVAC plants. *Journal of Process Control* 90, 1–17.
- Lu, Q., L. D. González, R. Kumar, and V. M. Zavala (2021). Bayesian optimization with reference models: A case study in MPC for HVAC plants. *Computers & Chemical Engineering* 154.
- Raissi, M., P. Perdikaris, and G. E. Karniadakis (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* 378, 686–707.
- Shariari, B., K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE* 104(1), 148–175.
- Snoek, J., O. Rippel, K. Swersky, R. Kiros, N. Satish, N. Sundaram, M. Patwary, M. Prabhat, and R. Adams (2015). Scalable Bayesian optimization using deep neural networks. In *International Conference on Machine Learning*, pp. 2171–2180. PMLR.
- Yamashita, A., A. Zanin, and D. Odloak (2016). Tuning of model predictive control with multi-objective optimization. *Brazilian Journal of Chemical Engineering* 33(2), 333–346.
- Zhang, J., S. D. Petersen, T. Radivojevic, A. Ramirez, A. Pérez-Manríquez, E. Abeliuk, B. J. Sánchez, Z. Costello, Y. Chen, M. J. Fero, et al. (2020). Combining mechanistic and machine learning models for predictive engineering and optimization of tryptophan metabolism. *Nature Communications* 11(1), 1–13.