

Neural networks smart grid based optimisation for expensive functions

Alain UWADUKUNZE^{1,2}, Xavier BOMBOIS³, Marion GILSON², Marie ALBISSER¹

Abstract—Bayesian optimisation is an emerging machine learning technique known to be efficient especially for optimising functions which are expensive to evaluate. In Bayesian optimisation, a Gaussian process model of the unknown function is identified based on available data. Its estimate of the unknown function and the associated uncertainties are used to build a so-called acquisition function which does a trade-off between exploitation and exploration. The latter is then iteratively maximised to find candidates which are promising to be close to the optimum. In this paper, an alternative version of Bayesian optimisation, where the Gaussian process model is replaced by a neural network model, is proposed. As shown in the numerical illustration of this paper, this alternative version will require less computation time when facing optimisation problems with initially large data sets. Since neural networks do not naturally provide an information about the quality of the estimates, a different strategy for the exploration objective of our approach is proposed.

I. INTRODUCTION

In many applications such as engineering design or control, one seeks to determine the maximum X_{opt} of an unknown static function $f(X)$ of a vector X of design variables. Even though $f(X)$ is unknown, this function can be evaluated for any given X (in a certain range). However, it is supposed that evaluating $f(X)$ for a given X is costly (in time and/or resources). Consequently, a gradient-based optimisation with a numerical evaluation of the gradient is not the most appropriate approach to determine X_{opt} . Instead, the literature proposes different techniques that are based on a smart gridding of the space of X . For each grid point X , the function is evaluated and a procedure determines the next grid point in a smart way. To achieve this, a model of the function $f(X)$ is identified based on the previous grid points $\{X^n, f(X^n)\}_{n=1}^k$ and this model is used to determine a point X^{k+1} (the next grid point) that is a promising candidate for (being close to) X_{opt} . A popular version of this smart grid approach is the so-called Bayesian optimisation (BO) [1] where the model of the function $f(X)$ takes the form of a Gaussian process (GP) [2]. In this paper, we will be comparing the classic BO to an alternative version, where the model of the function is a neural network (NN).

The problem described in the previous paragraph is quite classical in data-driven control applications where a controller achieving the highest level of performance is designed without requiring a model of the to-be-controlled

system. Instead, an optimal controller is determined by testing different controllers (with different parameters X) on the system. It is clear that, for many real-life systems, the evaluation of a candidate can be quite costly and the to-be-tested controllers should therefore be chosen in a smart way. Bayesian optimisation is generally performed for this purpose: a model of the performance $f(X)$ of the closed loop made up of the system and the controller with parameters X is identified based on the previous experiments and this model is used to determine the parameters of the controller that has to be tested in the next experiment. Examples of application of BO in control can be found in [3], [4], [5] and [6]. Bayesian optimisation is also performed in many other applications such as in *Environmental Monitoring* as suggested in [7].

The model of the unknown function $f(X)$ is a crucial component of such smart-grid based optimisation procedures. As mentioned above, GPs are generally used for this purpose. A GP is a non-parametric identification method that allows to deduce an estimate of $f(X)$ for any point X based on data points $\{X^n, f(X^n)\}_{n=1}^k$. Besides this estimate, the GP also provides an uncertainty region indicating the confidence we can have in the estimate. This uncertainty region is deduced in a Bayesian framework and its reliability of course depends on the chosen prior assumptions and kernels used in this Bayesian framework. This means that it is not guaranteed that the unknown true function lies in the uncertainty region. However, it will be generally observed that the uncertainty will be large for points X that are far away from the data points with which the GP has been identified.

The GP model of the function $f(X)$ (and its uncertainty) is then used to determine the next point X^{k+1} for which $f(X)$ must be evaluated. This is generally done by maximising a so-called *acquisition function* $A(X)$. The acquisition function $A(X)$ is a function of the GP model and of its uncertainty. Its maximum (*i.e.*, the next grid point X^{k+1}) will be a point for which the GP model of $f(X)$ is large (exploitation) and/or for which the uncertainty of the GP is high (exploration). The acquisition function achieves the trade-off between these two objectives.

In recent studies, the use of other types of models, in procedures similar to BO, has been explored. For example in [8], Inverse Distance Weighting interpolation and Radial Basis functions are successfully used as models of the to be optimised function. The goal in this paper, will be to show that a neural network (NN) model can also be efficiently used in a smart-grid based optimisation procedure. Neural networks are increasing in popularity and being applied in many fields such as system identification [9], [10]. Moreover,

¹ French-German Research Institute, ISL, 5 rue du Général Cassagnou, 68300 Saint-Louis, France

² Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France

³ Laboratoire Ampère, Ecole Centrale de Lyon, Université de Lyon, 36 avenue Guy de Collongue, Ecully, France and Centre National de la Recherche Scientifique (CNRS), France

they are known to require less training time than regular GP when the number of training samples increases. This may prove to be useful when BO needs to be performed with an initially large dataset. It is clear that, unlike a GP model, a NN model is generally not accompanied by an uncertainty region. However, it will be shown that efficient acquisition functions can also be built when the model of $f(X)$ is a NN model by using the distance to observed samples for the exploration. It will also be illustrated, through different case studies, that using a NN in a smart grid based procedure can provide similar optimisation performances to regular BO but presents the interesting advantage of requiring less computational time than regular BO when the initial dataset available is large.

The rest of the paper is structured as follows: First, the BO algorithm is presented in section III-A, then the proposed approach is presented in section III-B. Finally, the proposed method is illustrated through an application in section IV.

II. PROBLEM STATEMENT

This paper focuses on optimisation problems of the form:

$$X_{opt} = \arg \max_X f(X), \quad (1)$$

where $X \in \mathbb{R}^N$ with N the dimension of the inputs and $f(X) \in \mathbb{R}$. The function f is considered to be unknown and expensive to evaluate (*i.e.* it requires a lot of resources or time to obtain the value of $f(X)$ for a given X).

For the optimisation problem (1), this paper has the following objective: given a dataset, $D^k : \{X^n, f(X^n)\}_{n=1}^k$, obtained from an unknown function f , use a neural network as a model of the unknown function f in a smart grid based optimisation procedure such as BO, to search which promising candidates to evaluate to get close to X_{opt} .

III. METHODOLOGY

The proposed smart grid optimisation is inspired of BO which is briefly recalled here.

A. Bayesian optimisation

Bayesian optimisation aims at finding an optimum X_{opt} of the problem (1). The optimum is searched by iteratively maximising a so called *acquisition function* $A(X)$. The maximisation of $A(X)$ yields a promising candidate to evaluate on f in order to get close to X_{opt} . The acquisition function balances the search between exploitation (searching for the place where a model's estimate of f is large) and exploration (searching for the place where the uncertainties are the largest). Therefore, a model of the unknown function f and its uncertainty are needed to build the acquisition function. In this sense, all the available samples from f , $D^k : \{X^n, f(X^n)\}_{n=1}^k$, with k the number of samples, are used to construct a probabilistic model of $f(X)$ capable of providing uncertainties. The most commonly used probabilistic models are Gaussian processes (GP) [2] which are non-parametric regression models. A GP assumes a prior Gaussian distribution over an unknown function f defined by a prior mean $\mu_p(X)$ and Kernel matrix $K(X, X')$. Kernels are covariance

matrices which measure the similarity between points using distance metrics. An example of a frequently used kernel is the *Radial Basis Function* (RBF), defined as follows for two points X_a and X_b :

$$K(X_a, X_b) = \exp\left(-\frac{d(X_a, X_b)^2}{2l^2}\right), \quad (2)$$

where d is the Euclidean distance and l a length scale parameter. When samples from f , $D^k : \{X^n, f(X^n)\}_{n=1}^k$, become available, the GP model is updated through Bayesian inference and provides a Gaussian distribution $\mathcal{N}(\mu(X), \sigma^2(X))$ over the unknown function f . This process is repeated at each iteration of BO as new samples become available. The mean $\mu(X)$ is the best estimation of $f(X)$ and is used for the exploitation. The variance $\sigma^2(X)$ provides the uncertainties on f for each point X and is used for the exploration.

Several acquisition functions for BO can be found in the literature, as suggested in [1]. Among them are: Probability of Improvement, Upper Confidence Bound and Expected Improvement. Since the proposed approach is inspired of the Upper Confidence Bound (UCB) acquisition function (see III-B for comparison), the paper focuses mainly on this technique. This acquisition function is perhaps the most straightforward to understand and is defined as follows:

$$UCB(X) = \beta\sigma(X) + \mu(X), \quad (3)$$

where $\sigma(X)$ and $\mu(X)$ are the standard deviation and predictive mean of the GP, respectively. The trade-off between exploitation and exploration is achieved by finding the point X for which σ and μ are both high. The parameter β can be used to control the degree of the trade-off. When β is chosen high more importance will be given to the exploration objective and when it is chosen low, more importance will be given to the exploitation objective. However, the value of β can also be chosen more appropriately to, *e.g.*, bound some cumulative regret as suggested in [11].

The BO algorithm using UCB as acquisition function to solve an optimisation problem (1), given an initial dataset D^k , is described in Algorithm 1. Note that if the dataset D^k is large, *i.e.* k is large, the first step of Algorithm 1 which involves identifying a new GP model at each iteration of BO may be time consuming.

Depending on the application or available data, one may want to use other models to represent the unknown functions. As a result, an alternative to BO where neural networks are used as models of the unknown function is proposed in the following section.

B. Proposed approach using neural networks

In this section, an alternative approach to BO is proposed. We refer to this method as Neural Network Smart Grid Optimisation (NNSGO). Similarly to BO, the idea behind this method is that the most promising point to evaluate on the real function f should be a trade-off between the exploitation and exploration objectives. The main difference between NNSGO and BO is that NN are employed as models of the functions rather than using probabilistic models such

as GPs. The great advantage of neural networks resides in their ability to handle different types of problems while requiring less computational time to learn large datasets (*i.e.* datasets D^k , where the number of observations k is large) than other machine learning algorithms, such as GPs. Many different types of NN exist. In this paper, multi layer perceptrons (MLP), which are NN with fully connected layers are used as models. They provide the relation between a set of inputs X , where $X \in \mathbb{R}^N$ with N the dimension of the inputs, associated to outputs $f(X)$, where $f(X) \in \mathbb{R}$, thanks to trained parameters known as weights and bias.

A MLP needs to be trained to provide the best estimates of an unknown function. The training consists in finding the best values of parameters that minimise an identification criterion commonly known as loss function for NNs. The loss function which is frequently used in regression to train the MLP is the Mean Squared Error (*MSE*) defined as:

$$MSE = \frac{1}{k} \sum_{n=1}^k \|\hat{f}(X^n) - f(X^n)\|^2, \quad (4)$$

where k is the number of samples, $f(X^n)$ the target output for X^n , $\hat{f}(X^n)$ the estimate of $f(X^n)$ by the NN and $\|A\| = \sqrt{A^T A}$ represents the Euclidean norm of a vector A . The most commonly used method to train a MLP is back-propagation which uses Gradient Descent [12] to iteratively update the parameters of the NN. Using this method requires the MLP to have differentiable activation functions such as *ReLU* or *tanh*.

1) *Exploration strategy for NNSGO*: Unlike GPs, regular NN do not come with a measure of their uncertainties. Therefore, the exploration when using NN has to be achieved in a different way to classic BO. In the proposed procedure, we suggest to perform the exploration with the objective of covering regions not already covered by the dataset (and therefore not already observed by the NN). The rationale behind this proposal is that when a point X to be predicted is far from the observed points, then it is likely that the prediction $\hat{f}(X)$ will differ from the real value $f(X)$. This is

Algorithm 1 Bayesian optimisation algorithm

for $j = 0, \dots, N_{it}$, with N_{it} the maximum number of iterations, **do**

1. Train GP with available set of samples. $D^{k+j} : \{X^n, f(X^n)\}_{n=1}^{k+j}$

2. Build acquisition function:

$$A_j(X) = \beta \sigma_j(X) + \mu_j(X)$$

3. The next point X^{k+j} to evaluate on f is

$$X^{k+j} = \arg \max_X A_j(X)$$

4. Evaluate X^{k+j} on f and add $(X^{k+j}, f(X^{k+j}))$ to the training samples D^k .

end for

return X with the highest $f(X)$ from $D^{k+N_{it}}$

also the case in GP, where the exploration is performed using the uncertainties which are based on the similarity between the evaluated points and the observed ones. Therefore, this idea is used to create a measure which is employed for the exploration objective of the proposed procedure. In that respect, $\delta_{min}(X)$ is defined as the Euclidean distance between a point X and the closest point to it from the points in the dataset $D_X^k = X | (X, f(X)) \in D^k$. It is defined as:

$$\delta_{min}(X) = \min_{X^n \in D_X^k} \|X - X^n\|. \quad (5)$$

This means that $\delta_{min}(X)$ is small when evaluating a point close to the observed points and increases when evaluating a point far from the observed ones. When $\delta_{min}(X)$ is low, it is likely that the estimate of $f(X)$, $\hat{f}(X)$ is close to the real value $f(X)$ and vice-versa.

2) *Illustration of the comparison between δ_{min} and the GP uncertainties*: The measure δ_{min} can be compared to the uncertainties provided by a GP model. To illustrate this, $k = 5000$ random samples are generated from a function $f(X)$ which corresponds to the 12 dimensional form of the Rosenbrock function defined as:

$$f(X) = \sum_{i=1}^{N-1} [100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2], \quad (6)$$

with $X = (x_1, x_2, \dots, x_N) \in \mathbb{R}^N$ and $N = 12$. The optimisation of this function using the proposed procedure will be performed in section IV. The goal of this illustration is to visualize how the NN performs and how the GP standard deviation behaves when the value of δ_{min} increases. For this purpose, a GP and a NN model of the 12 dimensional Rosenbrock function, are identified using the generated samples. In this example, training the GP model requires more computational time than training the NN model. This will be further discussed in section IV. A new dataset D_{new}^l containing $l = 500$ new samples $\notin D^k$ is generated. The GP and NN models are then used to predict the 500 unseen points from D_{new}^l . The following elements are then calculated:

- The standard deviation $\sigma(X)$ for $X \in D_{new}^l$, note that this value is multiplied by 4 for better visualization in Fig. 1.
- The $\delta_{min}(X)$ for $X \in D_{new}^l$ and computed with (5) where D_X^k is the set containing the $k = 5000$ initial points.
- The error of the NN for each point $X \in D_{new}^l$ determined as the absolute value of the difference between the estimate of the Rosenbrock function by the NN $\hat{f}(X)$ and the actual value $f(X)$: $error = |\hat{f}(X) - f(X)|$.

Fig. 1 presents the GP standard deviation (σ) and the NN error w.r.t to δ_{min} for different tested points. On this figure, the value on the vertical axis of each orange dot represents the error on the NN prediction ($|\hat{f}(X) - f(X)|$) of an unseen tested point X and the corresponding value on the horizontal axis, represents the value of δ_{min} for that point.




Fig. 1: Value of GP σ (represented by the blue stars) and NN error (represented by the orange dots) w.r.t δ_{min} (horizontal axis).

The value on the vertical axis of each blue star represents the GP standard deviation σ of an unseen tested point X and the corresponding value on the horizontal axis, represents the value of δ_{min} for that point. We can observe that the points X for which $\delta_{min}(X)$ is low (*i.e.* points in regions well covered by the dataset) are relatively well predicted by the NN. The GP standard deviation $\sigma(X)$ is also low for these points. However, we see that when $\delta_{min}(X)$ increases, more and more points tend to be poorly predicted by the NN as the error increases. The GP standard deviation is also greater for points where $\delta_{min}(X)$ is large. Therefore, by using the measure $\delta_{min}(X)$ for the exploration objective of our procedure, we will test in regions where the model is likely to perform poorly. This is essential, because if the optimum is in those regions and therefore poorly predicted, without the exploration it may never be tested. This will also allow to improve the quality of the model in areas not already covered by the dataset.

3) *Acquisition function for the NNSGO*: The measure δ_{min} will therefore be used to construct the acquisition function for NNSGO. For this purpose, the problem (1) is still considered with initial samples $D^k : \{X^n, f(X^n)\}_{n=1}^k$ drawn from an unknown function f . The function f is modelled by a NN model $\hat{f}(X)$ which is trained with the initial samples D^k . We then choose the following expression for the acquisition function $A(X)$ in the NNSGO procedure:

$$A(X) = \beta \delta_{min}(X) + \hat{f}(X). \quad (7)$$

Compared to the UCB (3), $\hat{f}(X)$ plays the role of the GP's μ and serves for the exploitation purpose, δ_{min} serves the exploration purpose just like σ in (3) and β is a trade-off hyper-parameter which balances the search between exploration and exploitation. First of all, β must be chosen in a way that both the exploration and exploitation terms have the same order of magnitude. Afterwards, the value of β can be increased or decreased to permit more exploration or more exploitation, respectively. In our procedure, the first

$N_{it} - 1$ points to evaluate, with N_{it} the maximum number of iterations, are chosen by achieving the trade-off between the exploration and exploitation objectives (*i.e.* by maximisation of the acquisition function (7)). Since the real function can only be evaluated once after the final iteration, performing exploration in the objective of improving the model is no longer useful. The last point to evaluate is thus, chosen solely through the exploitation objective *i.e.* by finding the point X for which the NN estimate of f is the largest. The main steps of the NNSGO are described in algorithm 2.

4) *Distance metrics in higher dimensions*: The NNSGO requires the use of distance metrics. However, in higher dimensions when measuring similarities between points, they tend to provide poor contrasts between farthest and nearest neighbours to an evaluated point [13]. This makes the proposed procedure less efficient in higher dimensions. Nevertheless, two solutions can be considered in order to overcome this issue:

- Use fractional distance metrics which give better contrast between farthest and nearest neighbours to an evaluated point, as defined in [13], instead of the Euclidean distance.
- Increase the term δ_{min} to the power α (with $\alpha > 1$) to increase the contrast. Note that this approach will only be efficient if the evaluated δ_{min} is greater than 1.

IV. APPLICATIONS

This section follows on from section III-B.2: the objective is to perform the optimisation of the 12 dimensional form of the Rosenbrock function $f(X)$ (6). In this case study, the proposed NNSGO will be compared to BO. Before performing the procedures, a preliminary study is carried out to find the structures of the NN and of the GP models that best fit

Algorithm 2 NN Smart grid optimisation algorithm

for $j = 0, \dots, N_{it} - 1$ **do**

1. Train neural network model $\hat{f}(X)$ with available set of samples $D^{k+j} : \{X^n, f(X^n)\}_{n=1}^{k+j}$

2. Build acquisition function $A_j(X) = \beta \delta_{min}(X) + \hat{f}(X)$

3. The next point X^{k+j} to evaluate on f is

$$X^{k+j} = \arg \max_X A_j(X)$$

4. Evaluate X^{k+j} on f and add $(X^{k+j}, f(X^{k+j}))$ to the training samples D^{k+j} .

end for

Train neural network model $\hat{f}(X)$ with available set of samples $D^{k+N_{it}-1}$

The final point to evaluate on f is:

$$X^{k+N_{it}} = \arg \max_X \hat{f}(X)$$

Evaluate $X^{k+N_{it}}$ on f and add $(X^{k+N_{it}}, f(X^{k+N_{it}}))$ to the training samples $D^{k+N_{it}}$.

return X with the highest $f(X)$ from $D^{k+N_{it}}$

the data using *e.g* a grid search approach [14]. The NN and GP models are built using *Scikit-learn* [15], a famous Python machine learning library which offers the possibility to build simple MLP regression models and GPs. For more advanced NN models, tools such as *PyTorch* or *Keras* could also be used. The optimisation of different acquisition functions is achieved using the Limited-memory BFGS [16] method in its *Scipy* [17] implementation. Simulations are performed on a PC equipped with a 12th Gen Intel CORE i7-12700H processor with 14 cores and 64 GB of RAM.

The Rosenbrock function has one global minimum at $X = (1, 1, \dots, 1)$ for $f(X) = 0$. The optimisation is performed here in the search domain

$$\mathcal{X} = \{X = (x_1, \dots, x_{12})^T \mid -2.048 < x_i < 2.048 (i = 1, \dots, 12)\} \text{ where } X \in \mathbb{R}^{12}.$$

For the purpose of this study, it is considered that the function is unknown and expensive (in time and resources) to evaluate. However, we have access to a dataset containing 5000 samples (pairs $\{X^n, f(X^n)\}_{n=1}^k$ with $k = 5000$) and we have the possibility to test 100 additional points X . The dataset is generated randomly in a corner of the search domain which does not include the global minimum to motivate the use of a smart grid based approach.

The goal of this section will be to find the optimum X_{opt} of the following optimisation problem:

$$X_{opt} = \arg \max_{X \in \mathcal{X}} -f(X). \quad (17)$$

To solve this optimisation problem, the NNSGO is performed using algorithm 2 with the following considerations:

- The number of iterations is set to 100.
- The NN model used is a multi layer perceptron with 1 hidden layer with 256 neurons and a *ReLU* activation function on each neuron.
- The maximum of the acquisition function provides the next point X to evaluate and is given as:

$$A(X) = -\hat{f}(X) + \beta \delta_{min}(X), \quad (18)$$

where $\hat{f}(X)$, is the multi layer perceptron's estimate of $f(X)$ and $\beta = 1$.

The NNSGO is compared to BO which is performed using algorithm 1 with the following parameters:

- The number of iterations is set to 100.
- A GP with a Matern kernel [2] is used.
- The maximum of the acquisition function which provides the next controller to evaluate corresponds to the UCB (3) where $\mu(X)$ is the GP estimate of $-f(X)$, $\sigma(X)$ the standard deviation provided by the GP model and $\beta = 2$.

The values of β in algorithm 1 (BO) and algorithm 2 (NNSGO) are chosen to have the same order of magnitude between the exploration and exploitation terms.

Table I presents the values of f for the optimum points found using the NNSGO procedure (noted X_{NNSGO}) and BO (noted X_{BO}), as well as the time to run 100 iterations of both procedures. As shown in this table, both methods perform similarly in terms of optimums found. In the initial

TABLE I: Optimisation results

Method used	Optimisation time	Value of optimum
NNSGO	40 seconds	$f(X_{NNSGO}) = 40$
BO	1 hour	$f(X_{BO}) = 38$

dataset, the point X_D , for which the value of the Rosenbrock function is the lowest is for the value $f(X_D) = 226$. By increasing the dataset with 100 additional points using the two smart grid based approaches (BO and NNSGO), we are able to find optimum points X for which the value of $f(X)$ is 84% smaller than the initially smallest value in the dataset ($f(X_D) = 226$).

Even if the optimums found with both approaches (BO and NNSGO) have close values of $f(X)$, the NNSGO has a clear advantage in terms of the computational time needed to perform 100 iterations in this case study. This is due to the fact that neural networks scale better in training time than regular GPs when the size of the dataset increases. In this particular case, the NN takes 8 seconds to train whereas the GPs take 30 seconds. Furthermore, after each iteration of the NNSGO, the previously trained model parameters are re-used as an initialization for the new model. This further decreases the training time of the updated NN models. This is achieved by using the *warm start* option in *scikit-learn* multi layer perceptron regressor implementation. Without using this option, the optimisation when using the NNSGO procedure takes 12 minutes, which is still 6 times less than when using the regular BO to perform 100 iterations. (Note that the *warm-start* option is not available for GP regression models in *scikit-learn*).

Other case studies are available in sections IV.A and IV.B of [18] (which is the paper that we submitted for review). The paper [18] is available at <https://hal.science/hal-04052060/>.

V. CONCLUSIONS AND PERSPECTIVES

In this study, an alternative method to Bayesian optimisation for the optimisation of costly unknown functions is proposed. It is based on the use of neural networks instead of Gaussian processes. This method uses the distance between points, instead of the usual uncertainties provided by Gaussian processes, to search for a potential optimum to be tested next on the unknown function. The proposed procedure provides competitive results to regular Bayesian optimisation. It is also illustrated that when the initial dataset contains a large number of samples, the optimisation procedure using the proposed approach will require less time to perform a certain number of iterations than the classic Bayesian optimisation algorithm. The next step will be to improve this approach by extending it to optimisation problems under constraints.

REFERENCES

- [1] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.
- [2] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer school on machine learning*, pp. 63–71, Springer, 2003.

- [3] X. Bombois and M. Forgione, "Control design via Bayesian Optimization with safety constraints," *Conference on Control Technology and Applications*, 2022.
- [4] S. Bansal, R. Calandra, T. Xiao, S. Levine, and C. J. Tomlin, "Goal-driven dynamics learning via bayesian optimization," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 5168–5173, 2017.
- [5] S. Kuindersma, R. Grupen, and A. Barto, "Variational bayesian optimization for runtime risk-sensitive control," *Robotics: Science and Systems VIII*, pp. 201–208, 2012.
- [6] R. Guzman, R. Oliveira, and F. Ramos, "Heteroscedastic bayesian optimisation for stochastic model predictive control," *IEEE Robotics and Automation Letters*, vol. 6, no. 1, pp. 56–63, 2020.
- [7] R. Marchant and F. Ramos, "Bayesian optimisation for intelligent environmental monitoring," in *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 2242–2249, IEEE, 2012.
- [8] A. Bemporad, "Global optimization via inverse distance weighting and radial basis functions," *Computational Optimization and Applications*, vol. 77, no. 2, pp. 571–595, 2020.
- [9] C. Andersson, A. H. Ribeiro, K. Tiels, N. Wahlström, and T. B. Schön, "Deep convolutional networks in system identification," in *2019 IEEE 58th conference on decision and control (CDC)*, pp. 3670–3676, IEEE, 2019.
- [10] M. Forgione and D. Piga, "dynonet: A neural network architecture for learning dynamical systems," *International Journal of Adaptive Control and Signal Processing*, vol. 35, no. 4, pp. 612–626, 2021.
- [11] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," *arXiv preprint arXiv:0912.3995*, 2009.
- [12] S. Ruder, "An overview of gradient descent optimization algorithms," *arXiv preprint arXiv:1609.04747*, 2016.
- [13] C. C. Aggarwal, A. Hinneburg, and D. A. Keim, "On the surprising behavior of distance metrics in high dimensional space," in *International conference on database theory*, pp. 420–434, Springer, 2001.
- [14] P. Liashchynskiy and P. Liashchynskiy, "Grid search, random search, genetic algorithm: a big comparison for nas," *arXiv preprint arXiv:1912.06059*, 2019.
- [15] L. Buitinck, G. Louppe, M. Blondel, F. Pedregosa, A. Mueller, O. Grisel, V. Niculae, P. Prettenhofer, A. Gramfort, J. Grobler, R. Layton, J. VanderPlas, A. Joly, B. Holt, and G. Varoquaux, "API design for machine learning software: experiences from the scikit-learn project," in *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, pp. 108–122, 2013.
- [16] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, "A limited memory algorithm for bound constrained optimization," *SIAM Journal on scientific computing*, vol. 16, no. 5, pp. 1190–1208, 1995.
- [17] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [18] A. Uwadukunze, X. Bombois, M. Gilson, and M. Albisser, "Neural networks smart grid based optimisation for expensive functions," 2023. Available on HAL repository at <https://hal.science/hal-04052060/>.