Hospital Simulation Model Optimisation with a Random ReLU Expansion Surrogate Model

Laurens Bliek*
l.bliek@tue.nl
Eindhoven University of Technology
Eindhoven, Netherlands

Arthur Guijt*
arthur.guijt@cwi.nl
Centrum Wiskunde & Informatica
Amsterdam, Netherlands

Rickard Karlsson*
rickkarl@student.chalmers.se
Chalmers University of Technology
Gothenburg, Sweden

ABSTRACT

The industrial challenge of the GECCO 2021 conference is an expensive optimisation problem, where the parameters of a hospital simulation model need to be tuned to optimality. We show how a surrogate-based optimisation framework, with a random ReLU expansion as the surrogate model, outperforms other methods such as Bayesian optimisation, Hyperopt, and random search on this problem.

CCS CONCEPTS

• **Mathematics of computing → Mathematical optimization**; • **Computing methodologies →** *Active learning settings.*

KEYWORDS

surrogate models, expensive optimisation, simulation optimisation

ACM Reference Format:
Laurens Bliek, Arthur Guijt, and Rickard Karlsson. 2021. Hospital Simulation Model Optimisation with a Random ReLU Expansion Surrogate Model. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3463279

1 INTRODUCTION

The industrial challenge of the GECCO 2021 conference is a relevant optimisation problem from healthcare that contains various challenging properties such as noise and an expensive objective function [9]. The objective function is based on a simulator that is used for resource planning tasks in hospitals under the specific situation of the COVID-19 pandemic. The goal is to tune the parameters of the simulator such that it can be used to provide accurate predictions.

Surrogate-based optimisation algorithms are designed to deal with expensive objective functions by approximating the objective with a surrogate model and then using this model to guide the optimisation procedure. In this work we showcase the surrogate algorithm that we use to tackle this challenge. In the remainder of this section, we introduce the optimisation problem that needs to be solved in this challenge, and we give an introduction on surrogate-based optimisation and the specific surrogate model that we use in our approach. Section 2 describes the results that we have achieved on this problem, as well as a comparison with related approaches. We conclude this work in Section 3.

1.1 Hospital Simulation Model Optimisation

The problem under consideration is to find the optimal set of parameters for a Hospital Simulation model, *BaBSim.Hospital*. This simulator is used for resource planning and allocation at hospitals by simulating different discrete events, such as an increase in the intake of patients or other scenarios during the COVID-19 pandemic [1, 2].

The simulator has a relatively long running time, on average 86 seconds on our machine, and has 29 parameters that need to be configured. Each parameter is real-valued and has a recommended range of values to limit the search space. More formally, we define an output from the simulator $f(x)$ that represents the objective score of a proposed set of parameters x. Then, the goal is to find x^* with the lowest objective score possible. This equates to solving the following minimisation problem

$$\begin{aligned} &\underset{x}{\text{argmin}} && f(x) \\ &\text{subject to} && l_k \le x_k \le u_k \quad k = 1, \ldots, d \end{aligned}$$

where l_k and u_k denote the lower and upper bounds for the k-th variable, respectively, and $d = 29$ is the number of variables.

On top of that, however, we are also limited to a fixed number of 200 evaluations of $f(x)$. This puts more emphasis on an efficient search through the solution space. It also means that we might not expect to solve the problem optimally. To make things more complicated, the simulator is non-deterministic, and the observed evaluation of $f(x)$ is therefore perturbed by random noise.

1.2 Surrogate-based optimisation

To deal with expensive objective functions using a low number of function evaluations, we make use of a surrogate-based optimisation framework. In this framework, we perform three steps at every iteration i:

(1) Evaluate objective function f at candidate point $x^{(i)}$.
(2) Approximate f with a surrogate model g.
(3) Use g to suggest a new point $x^{(i+1)}$.

This is the approach taken by, for example, Bayesian optimisation algorithms [11]. Usually, the first r candidate points are chosen randomly, which means step 3 is skipped for $i = 1, \ldots, r$. For step 2, typical choices of the surrogate model g are Gaussian processes [11], Parzen estimators [3] or random feature expansions [5].

*All three authors contributed equally to this research.

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07.
https://doi.org/10.1145/3449726.3463279

Various methods exist to use the surrogate model to suggest a new point $x^{(i+1)}$. In this work, we use a continuous optimiser named L-BFGS [12] to find a local minimum of g, and then add a random perturbation, changing step 3 into:

$$x^{(i+1)} = \underset{x}{\operatorname{argmin}}\, g(x) + \delta, \tag{1}$$

where δ_k, the element of vector δ corresponding to the k-th variable, is a normally distributed perturbation parameter $\delta_k \sim N(0, \sigma_k)$. For the standard deviation we chose $\sigma_k = 0.1(u_k - l_k)/\sqrt{d}$. The goal of the random perturbation is to avoid getting stuck in local optima and to stimulate exploration of the search space.

1.3 Random ReLU Expansion Surrogate Model

The surrogate model used in this work is a random ReLU expansion (RRE), where ReLU stands for rectified linear unit - a common basis function used in the machine learning community. A RRE is a specific case of random feature expansion

$$g(x) = \sum_{k=1}^{D} c_k \varphi(w_k^T x + b_k), \tag{2}$$

for which the approximation capabilities are well understood [8], with φ equal to a ReLU. The ReLU is defined as

$$\varphi(z) = \begin{cases} z, & z > 0, \\ 0, & z \leq 0, \end{cases} \tag{3}$$

which causes the surrogate model g to be piece-wise linear. Other versions of the random ReLU expansion were introduced in [4] and [6], but unlike those approaches we do not enforce any convexity or integer constraints in this work. The model parameters w_k and b_k are chosen randomly but stay fixed for the whole duration of the algorithm, as is common in random feature expansions [5, 8]. The only model parameters that are trained are the parameters c_k, for which we use the recursive least squares algorithm [10] with a regularisation factor of 10^{-8}. Finally, the total number of basis functions is set to $D = 1000$.

2 RESULTS

Using a benchmark suite for expensive optimisation problems [7], we compared our approach to two surrogate algorithms: Bayesian optimisation[1] and HyperOpt [3], also known as TPE. Default values were chosen for these algorithms. For these two algorithms, as well as for our RRE approach, we started with $r = 2d = 58$ random iterations from a uniform distribution, as common choices for this number are $r = d$ and $r = 2d$, with d the number of variables. We also compared with a pure random search, which is equivalent to using any surrogate method with $r = 200$, as the total number of function evaluations equals 200.

Table 1 shows the best objective function value found by each method after 200 expensive evaluations of the Hospital Simulation objective function, averaged over five runs. Our RRE approach achieves the best objective value out of all methods, namely 16.29, however its variance is quite high and even comparable with random search. Still, we choose to further fine-tune our approach while disregarding other approaches, and look for ways to reduce the variance. We speculate that the large variance comes from the large number of random evaluations r, which limits the number of points that are suggested by the surrogate model. Indeed, when lowering r to $r = d = 29$, the variance is reduced, and the average best objective value is reduced to 14.81 as well.

[1] We use the implementation from https://github.com/fmfn/BayesianOptimization, which uses a Gaussian process with a Matérn 5/2 kernel.

Table 1: Lowest objective value for different surrogate algorithms after 200 iterations, averaged over five runs.

Method	Result
Random search	17.65 (±2.17)
HyperOpt (TPE)	16.81 (±1.01)
Bayesian optimisation	16.78 (±1.32)
RRE ($r = 58$)	16.29 (±2.16)
RRE ($r = 29$)	14.81 (±0.69)

3 CONCLUSION

We have applied a benchmark suite of approaches to investigate the performance of surrogate methods on the parameter estimation of a Hospital Simulation model. We found that using random ReLU expansions as a surrogate model provides the best performance in this scenario. However, with 58 initial random samples this approach suffers from high variability. Halving the number of random samples increases performance notably and results in a reduction in variability. In the future, we will investigate other parts of our approach, such as the exploration strategy, in order to achieve even better results on this challenging problem.

REFERENCES

[1] Thomas Bartz-Beielstein, Eva Bartz, Frederik Rehbach, and Olaf Mersmann. 2020. Optimization of High-dimensional Simulation Models Using Synthetic Data. arXiv:2009.02781 [stat.AP]
[2] Thomas Bartz-Beielstein, Frederik Rehbach, Olaf Mersmann, and Eva Bartz. 2020. Hospital Capacity Planning Using Discrete Event Simulation Under Special Consideration of the COVID-19 Pandemic. arXiv:2012.07188 [stat.AP]
[3] James Bergstra, Daniel Yamins, and David Cox. 2013. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International conference on machine learning*. 115–123.
[4] Laurens Bliek, Michel Verhaegen, and Sander Wahls. 2017. Online function minimization with convex random ReLU expansions. In *MLSP*. IEEE, 1–6.
[5] Laurens Bliek, Hans R. G. W. Verstraete, Michel Verhaegen, and Sander Wahls. 2018. Online Optimization With Costly and Noisy Measurements Using Random Fourier Expansions. *IEEE Transactions on Neural Networks and Learning Systems* 29, 1 (2018), 167–182.
[6] Laurens Bliek, Sicco Verwer, and Mathijs de Weerdt. 2020. Black-box Mixed-Variable Optimisation using a Surrogate Model that Satisfies Integer Constraints. *arXiv preprint arXiv:2006.04508* (2020).
[7] Arthur Guijt, Laurens Bliek, Rickard Karlsson, Sicco Verwer, and Mathijs de Weerdt. 2021. EXPObench: An EXPensive Optimization benchmark library. https://github.com/AlgTUDelft/ExpensiveOptimBenchmark.
[8] Ali Rahimi and Benjamin Recht. 2008. Uniform approximation of functions with random bases. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*. IEEE, 555–561.
[9] Margarita Rebolledo, Frederik Rehbach, Sowmya Chandrasekaran, and Thomas Bartz-Beielstein. 2021. Optimization of a simulation model for a capacity and resource planning task for hospitals under special consideration of the COVID-19 pandemic. https://www.th-koeln.de/informatik-und-ingenieurwissenschaften/gecco-2021-industrial-challenge-call-for-participation_82086.php.
[10] Ali H. Sayed and Thomas Kailath. 1998. Recursive least-squares adaptive filters. *The Digital Signal Processing Handbook* 21, 1 (1998).
[11] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* 104 (2016), 148–175.
[12] Stephen Wright and Jorge Nocedal. 1999. Numerical optimization. *Springer Science* 35 (1999), 67–68.