lajss.org©
Latin American Journal of Solids and Structures
www.lajss.org

Noise Handling in Kriging-Based Optimization Algorithms Applied to Sequential Decision Problems in Infrastructure Planning

Cibelle Dias de Carvalho Dantas Maia[a*] https://orcid.org/0000-0002-2410-6964, Rafael Holdorf Lopez[b] https://orcid.org/0000-0001-9037-0176, Tiago Luiz Schmitz[c] https://orcid.org/0009-0001-8389-4745

[a] Federal University of Campina Grande, 1770 Jairo Vieira Feitosa St, Pombal, PB, 58840-000, Brazil. Email: cibelle.dias@professor.ufcg.edu.br
[b] Federal University of Santa Catarina, 205 João Pio Duarte Silva St, Florianópolis, SC, 88040-970, Brazil. Email: rafael.holdorf@ufsc.br
[c] Santa Catarina State University, 2007 Madre Benvenuta Ave, Florianópolis, SC, 88035-901, Brazil. Email: tiago.schmitz@udesc.br

* Corresponding author

Abstract
This paper presents a comparative study of stochastic Kriging-based optimization algorithms applied to a generic infrastructure planning problem, using direct policy search (DPS) as a heuristic approach. The evaluated problem is a sequential decision problem involving a generic infrastructure planning scenario under uncertainty, with performance dependent on system and cost model parameters. The focus is to compare the impact of heterogeneous noise treatment in an optimization framework, approaching three algorithms: Minimum Quantile criterion (MQ), stochastic Efficient Global Optimization (sEGO), and Expected Improvement with Reinterpolation (EIR). The MQ algorithm is the only one that does not address information about the noise variance of the stochastic parameter. The results demonstrate that the algorithms presented satisfactory performance, especially those with heterogeneous noise treatment, and show potential for solving complex infrastructure engineering planning problems under uncertainties in a DPS framework.

Keywords
Optimization problem, stochastic Kriging, heterogeneous noise, direct policy search, infrastructure planning

Graphical Abstract

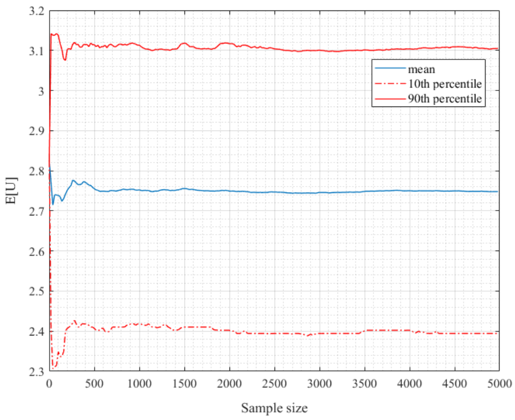

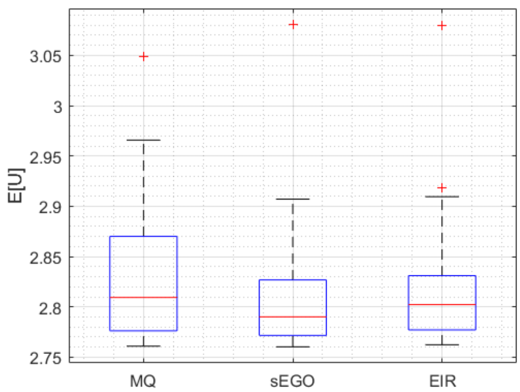

1 INTRODUCTION

Direct policy search is one of the solution frameworks for solving the general decision problem under uncertainty. This method has been applied to a wide range of infrastructure planning problems, including transportation (Golla et al., 2021), energy (Gupta et al., 2020), water resource management (Giuliani et al., 2016), and risk-based inspection (RBI) planning (Bismut and Straub, 2018; Luque and Straub, 2019).

In DPS with a heuristic, the parameters of a function mapping the system state to decisions are optimized rather than the decisions themselves. What results from optimization with DPS is therefore not a sequence of decisions, but a policy that one can operate (Quinn et al., 2017). A generic infrastructure planning problem consists of finding the heuristic parameters that maximize the expected total life-cycle utility under uncertainties (Bismut and Straub, 2019). Therefore, a global optimization method is necessary.

One strategy for solving stochastic optimization problems is through Kriging-based algorithms. The Kriging is a Gaussian based metamodel, that acts as an interpolating or regressor curve of support points that have information from objective function. This approach allows for predicting results without use of the expensive primary source (objective function) (Forrester and Keane, 2009). One Kriging structure, suitable for application to noisy problems, is a regression method known as Stochastic Kriging (SK) (Ankenman et al., 2010). One approach to improving this method, as used in this paper, is by replacing the standard point variance estimates with smoothed variance evaluations, as suggested by Kaminski (2015).

The primary motivation for using Kriging-based algorithms in optimization problems is to reduce the number of expensive fitness evaluations without degrading the quality of the obtained optimal solution, even with a small sample of stochastic parameters. Therefore, a great advantage of the Kriging surrogate is that it uses the mean squared error (MSE) to quantify the response surface's uncertainty (Van Beers and Kleijnen, 2003; Forrester and Keane, 2009).

In this article, we evaluate the effectiveness of three Kriging-based optimization algorithms applied to a heterogeneous noise DPS problem: Minimum Quantile criterion (MQ), stochastic Efficient Global Optimization (sEGO), and Expected Improvement with Reinterpolation (EIR).

The MQ algorithm, proposed by Cox and John (1992), balances exploration using metamodel variance information and exploitation using a percentile of the SK value. However, this method does not address the noise information in the stochastic DPS problem. In stochastic simulation, these algorithms that do not address information from function variance may not be very appropriate, as it ignores the noise in the observations, assuming that samples were taken with infinite precision.

A traditional approach to optimization through Kriging is the Efficient Global Optimization (EGO) method (Jones et al., 1998). An extension of EGO for stochastic problems that incorporates heterogeneous noise treatment is Stochastic Efficient Global Optimization (sEGO), proposed by Carraro et al. (2019).

In sEGO, the search for new points in the optimization process occurs in only one step, and the stochastic noise treatment information is represented directly in its formulation. Another heterogeneous noise treatment strategy involves a second iterative process. An example is the Expected Improvement with Reinterpolation (EIR) method proposed by Nascentes et al. (2019) for stochastic Kriging. In this approach, instead of modifying the algorithm structure for stochastic cases, the SK model is used to represent the problem, followed by using the SK predictions to create a new Kriging surrogate, now deterministic. With this noise-free second model, a classical deterministic optimization can be applied.

The goal of this paper is to compare the performance of those three stochastic Kriging-based optimization algorithms - MQ, sEGO, and EIR - in a DPS heuristic framework applied to a generic infrastructure planning problem. The objective is to find the parameters that maximize the expected total life-cycle utility under uncertainties.

This article is organized as follows: the second section presents the heuristic of the direct policy search applied to the generic planning problem. The third section details the Kriging-based optimization algorithms. The fourth section presents the analysis and results considering the problem in one-dimensional and three-dimensional cases. Lastly, conclusions are presented in the fifth section.

2 DIRECT POLICY SEARCH TO THE INFRASTRUCTURE PLANNING PROBLEM

Direct policy search (DPS) is a strategy for solving sequential decision problems that addresses heuristics as a search strategy for the global optimum within a reduced solution space. DPS is often chosen for its flexibility and intuitive principles (Bismut and Straub, 2019). The optimal strategy for the decision problem is:

(1)

where $\mathcal{S}$ is the space of all possible strategies, and $\mathbf{E}[y(x,\theta)]$, in our problem, is the expected total life-cycle utility associated with a strategy x, and , is the vector of stochastic parameters.

2.1 Infrastructure planning problem

We investigate a generic infrastructure planning problem, as described in Bismut and Straub (2019), which involves increasing the system's capacity in an optimized manner to meet demand at the lowest implementation cost. Therefore, each year t, the system's capacity a_t must cover the demand θ_t, which will increase over the discrete time interval [1, 2, 3, ..., T]. The initial system capacity a_1 is fixed by the operator and can be increased at any time for a cost. The cumulative update costs are given by:

$$U = y(\mathbf{x},\theta) = \sum_{t=1}^{T} U_{C,t}(\mathbf{x}) + U_{R,t}(\mathbf{x},\theta), \tag{2}$$

in which

(3)

(4)

$$U_{R,t}(\mathbf{x},\theta) = \Phi\left(\frac{a_t - \theta_t}{\alpha}\right) c_F \gamma_{t-1}, \tag{5}$$

where $c_a = 1$ is the upgrading cost factor, $c_F = 10$ is the penalty factor, Φ is the standard normal cumulative distribution of the capacity, $\Phi \sim N(a_t, \theta_t)$, $\alpha = 0.1$ is the tolerance, γ_t is the discount factor $\gamma_t = 1/(1+r)^t$ with $r = 0.02$ as used in Bismut and Straub (2019) and, lastly, θ_t is the system demand defined in Table 1.

The system incurs a penalty, $U_{R,t}$, when demand is not met within a certain margin. Thus, the expected total cost, U, will be given by the portion of the upgrade cost, $U_{C,t}$, plus the cost of the excess penalty, $U_{R,t}$. So, the demand growth is dependent on random quantities and the objective function. The expected value is given by:

$$E[y(\mathbf{x},\theta)] = E[U] = E\left[\sum_{t=1}^{T} U_{C,t}(\mathbf{x}) + U_{R,t}(\mathbf{x},\theta)\right]. \tag{6}$$

The parameters used for the DPS are provided in Table 1. Given the uncertain nature of demand, the initial demand θ_1 is modeled as a normal random variable. T is the design horizon, while Z_t denotes the noisy observation of demand at time, t.

Table 1 Model parameters.

Variable	Type	Mean	Std.Dev.
	Normal distr.		
	function		-
	Normal distr.		
	Normal distr.		
	Deterministic		-

After setting the initial capacity a_1, the system is subject to the demand of the first year θ_1, where Z_1 is the noisy observation of this demand. Then, the capacity for the next year (a_2) is calculated, and so on, up to the time horizon T. The capacity can only increase over time and must be restricted to six stages, $t = 0,1,2,3,4,5,6$.

2.2 Heuristics investigated

Figure 1 presents a pseudocode of the heuristic adopted for the infrastructure planning problem, as presented by Bismut and Straub (2019), to update the system's capacity. According to this approach, when the current observation Z_t

is within a specific margin of the available capacity a_t, the system's capacity is increased by at least Δa, where the size of this margin is determined by the factor k. To maximize the expected total life-cycle utility it is necessary to optimize the parameters a_1, Δa and k. For this purpose, three optimization algorithms based on Kriging - MQ, sEGO, and EIR — are used, and their results are compared.

```
in heuristic parameters {a_1 ≥ 0, Δa ≥ 0, k}
time horizon T, observations of demand Z_1, ..., Z_T
out vector a_t of capacities
t ← 1;
while t < T do
    if a_t − Z_t < k · a_t then
        a_th ← k · a_t + Z_t;
        a_{t+1} ← min (max(a_t + Δa, a_th), 6);
    end
    t ← t + 1;
end
return a_t
```

Figure 1: Pseudocode for the heuristic to upgrade the capacity based on the observed value of demand

3 KRIGING METAMODEL BASED OPTIMIZATION

The surrogate model is only an approximation of the true stochastic function we wish to optimize. However, the use of the Kriging metamodel can provide accurate predictions of complex landscapes and offers a credible estimate of the possible error in these predictions. So, in Kriging-based optimization algorithms, the error estimates make it possible to make tradeoffs between sampling where the current prediction is good (exploitation) and sampling where there is high uncertainty in the function predictor value (exploration), allowing searching the decision space efficiently (Jones et al., 1998).

Kriging-based optimization algorithms start by simulating a limited set of input combinations (referred to as initial sampling) and iteratively select new input combinations to simulate by evaluating an infill criterion (IC), which reflects information from Kriging. Those updates points are called Infill Points (IPs). The response surface is then updated sequentially with information obtained from the newly simulated IPs, improving the surrogate accuracy while at the same time seeking its global minimum. The procedure is repeated until the desired performance level is reached and the estimated optimum is returned (Rojas-Gonzalez and Van Nieuwenhuyse, 2020).

All evaluated algorithms used an updated version of stochastic kriging (SK) with smoothed variance, employing deterministic Kriging metamodel prediction to approximate the problem variance (Kaminski, 2015; Carraro et al., 2019). Both stochastic and deterministic Kriging use the same sampling plan $\chi = \left[x^{(1)}, x^{(2)}, \ldots, x^{(n_s)}\right]$, with n_s is the number of support points.

The prediction and variance of the SK at a given point x^+ is given by:

(7)

(8)

where $\delta = 1 - \mathbf{1}^{T}\left[\Sigma_Z + \hat{\Sigma}_\epsilon\right]^{-1}\hat{\sigma}_Z^2\,\mathbf{h}(x^+)$, $\hat{\mu}$ is the vector of estimated values for the mean at each design point, $\hat{\sigma}_z^2$ is the estimated value for the extrinsic variance associated with the implementation of the metamodel, $\mathbf{h}$ is the correlation vector between the point to be predicted and each of the design points, Σ_Z is the covariance matrix, $\hat{\Sigma}_\epsilon$ is the covariance matrix referring to the variance of the intrinsic noise, $\mathbf{1}$ is the unit vector, and $\bar{\boldsymbol{y}}$ is the approximate mean value of the stochastic function at each design point.

For notation convenience, the $\bar{y}$ estimates are grouped into the vector $\bar{y} = [\bar{y}(x^{(1)}), \bar{y}(x^{(2)}), \ldots, \bar{y}(x^{(n_s)})]^T$, and the estimatiors are given by:

(9)

where n_t is the number of samples of the stochastic parameter, i.e., the number of simulation replications taken at design points, and $\boldsymbol{\theta}_j$ is a random sample from a stochastic variable of the problem.

The correlation vector of x^+ with the n_s support points is given by:

(10)

where c_k is a gaussian kernel parameterized for predictor. The SK surrogate parameters $(\hat{\mu}, \hat{\sigma}_Z^2, c_k)$ are calculated via Maximum Likelihood Estimation (MLE):

, (11)

where

(12)

For the construction of the SK model, it is assumed that the observed data comes from a Gaussian process $f(x^{(i)}) = \mu + Z(x^{(i)})$, where μ is its constante mean, Z is a zero mean stationary gaussian process with variance σ_Z^2 and covariance $\Sigma_Z = \sigma_Z^2 \exp\left[-\sum_{k=1}^{n_x} c_k |x_k^{(i)} - x_k^{(j)}|^2\right]$.

By employing deterministic Kriging metamodel prediction to approximate the variance of the problem in stochastic Kriging, the covariance function is given by $\hat{\Sigma}_\epsilon = \mathrm{Diag}\left[\hat{V}(x^{(1)}), \ldots, \hat{V}(x^{(n_s)})\right]/n_t$, where $\hat{V}$ is an estimator of $V(x_i)$ given by a deterministic Kriging metamodel. Thus, the Kriging prediction for variance at a given point is:

(13)

where $\boldsymbol{h}$ is the correlation vector for variance, $\hat{\mu}_V$ are the mean trend of the variance on the Kriging metamodel, Σ_{Z_V} is the covariance matrix of all the support points of variance extrinsic noise, and $\bar{v}$ are the estimates of variance at each design point.

For notation convenience, the $\bar{v}$ estimates are grouped into the vector $\bar{v} = [\bar{V}(x^{(1)}), \bar{V}(x^{(2)}), \ldots, \bar{V}(x^{(n_s)})]^T$, and the estimatiors are given by:

(14)

The correlation vector of x^+ with the n_t support points is given by:

(15)

where c_{V_k} is a gaussian kernel parameterized for variance predictor. The Kriging surrogate parameters $(\hat{\mu}_V, \hat{\sigma}_{Z_V}^2, c_{V_k})$ are calculated via MLE:

, (16)

where

(17)

For the construction of the Kriging metamodel, it is assumed that the observed data comes from a Gaussian process $V(x^{(i)}) = \mu_V + Z_V(x^{(i)})$, where μ_V is its constante mean, Z_V is a zero mean stationary gaussian process with variance $\sigma^2_{Z_V}$ and covariance $\Sigma_{\mathbf{Z_V}} = \exp\left[-\sum_{k=1}^{n_x} c_{V_k}|x_k^{(i)} - x_k^{(j)}|^2\right]$.

These Kriging structures can be consulted in Kaminski (2015) and Carraro et al. (2019). The remainder of this section briefly explains the search and the replication strategy for each algorithm.

3.1 Minimum Quantile criterion - MQ

The Minimum Quantile (MQ) criterion, originally proposed by Cox and John (1992), aims to balance exploration and exploitation by selecting, as the next infill point (IP), the location that minimizes a chosen percentile of the predictor obtained via SK. In other words, it uses a weighted sum of the predicted mean, $\hat{h}$, and the predicted variance, $\hat{s}^2$ (Jalali et al., 2017). The quantile of the predicted value is given by:

(18)

where $\hat{y}$ is SK predictor defined in Equation 7, $\hat{s}^2$ is the standard deviation of the SK prediction defined in Equation 8, Φ is the cumulative probability density of the normal distribution and $\beta \in [0.5, 1]$ is a parameter that adjusts a quantile level, i.e. tunes the level of reliability wanted on the final result. In this context, $\beta = 0.5$, as adopted by Picheny et al. (2013), which means that the design points are evaluated based on the kriging mean predictor only, without taking into account the prediction variance at those points. So, the infill point in each iteration is:

$$\mathbf{x} = \arg\min_{x \in \mathcal{X}} \; MQ(x^+). \tag{19}$$

This method does not account for the noise variance associated with the stochastic parameter.

3.2 stochastic Efficient Global Optimization - sEGO

The sEGO algorithm proposed by Carraro et al. (2019) chooses the alternative with maximum augmented expected improvement (AEI) as the next infill point:

$$AEI(x^+) = E[max(y_{min} - \hat{y}, 0)] \left(1 - \frac{\hat{\sigma}^2_\epsilon(x^+)}{\sqrt{\hat{s}^2(x^+) + \hat{\sigma}^2_\epsilon(x^+)}}\right), \tag{20}$$

where y_{min} is the Kriging prediction at the current effective best solution, $\hat{\sigma}_\epsilon(x^+) = \hat{V}(x^+)/n_t$ is the variance of the noise intrinsic to the stochastic function, $\hat{y}$ is SK predictor defined in Equation 7 and $\hat{s}^2$ is SK variance defined in Equation 8.

They_{min} calculation assumes the point with minimum Kriging quantile, $q(x)$, among the simulated points, i.e., $q(x) = \hat{y} + \Phi^{-1}(\beta)\hat{s}(x)$ with $\beta = 0.84134$ (Jalali et al., 2017). In this case, the parameter β corresponds to a quantile approximately one standard deviation above the mean of the standard normal distribution, thus balancing exploration and exploitation. High values of β (i.e., near 1) penalize designs with high uncertainty, reflecting a more conservative approach. Hence, with a high β, the criterion is more likely to favor observation repetitions or clustering, to decrease locally the prediction variance. In contrast, with $\beta = 0.5$ in the MQ approach, the criterion can be expected to be more exploratory.

The first part of Equation 21 could be calculated as:

$$EI(\mathbf{x}^+) = \left(y_{min} - \hat{y}(\mathbf{x}^+)\right)\Phi\left(\frac{y_{min} - \hat{y}(\mathbf{x}^+)}{\hat{s}(\mathbf{x}^+)}\right) + \hat{s}(\mathbf{x}^+)\phi\left(\frac{y_{min} - \hat{y}(\mathbf{x}^+)}{\hat{s}(\mathbf{x}^+)}\right), \tag{21}$$

where Φ and ϕ are the cumulative distribution function and probability density function, respectively, and y_{min} is the smallest sampled value of y. So, in the optimization process, the next IP is found maximizing $\mathrm{AEI}(x^+)$, i.e., leads to the new point x^+ with the highest probability of improvement, either by sampling toward the optimum or improving the approximation of the metamodel.

3.3 Expected Improvement with Reinterpolation - EIR

Approached by Nascentes et al. (2019), the method proposes to use SK and deterministic Kriging together like a noise-handling strategy, instead of modifying the EI for cases stochastic.

The predictions by SK at the support points will be used to build a new model in Kriging, now deterministic, since the predictions will come free of intrinsic error. In other words, after training the SK model, its output serves as the training set for a new surrogate model that is now considered to be without noise. As this last model is noise-free, the classic EI, Equation 21, could be used as a metric to obtain new IPs (Forrester et al., 2006). So, the Kriging prediction for the deterministic case will be rewritten as:

(22)

$$\hat{\mu} = \frac{\mathbf{1}\Psi^{-1}\hat{\mathbf{y}}_r}{\mathbf{1}\Psi^{-1}\mathbf{1}}, \tag{23}$$

where $\hat{\mu}$ is the trend of the deterministic Kriging metamodel obtained with information about the response surface constructed through SK model $(\hat{y}_r)$, $\mathbf{h}$ is the correlation vector, Ψ is the covariance matrix of all support points, and $\bar{y}$ is the vector of the approximate mean value of the objective function at each design point.

And, as we can use the SK predictor itself, the re-interpolation will only need the value of the new spatial variance of the correlation between the support points.

$$\hat{\sigma}_{ri}^2 = \frac{1}{n}\left[(\bar{y} - 1\hat{\mu}_r)^T\hat{\Sigma}^{-1}\psi\hat{\Sigma}^{-1}(\bar{y} - 1\hat{\mu}_r)\right]. \tag{24}$$

4 COMPARISON OF RESULTS

In this section, we investigated the performance of the three stochastic Kriging-based algorithms optimization, which are MQ, sEGO, and EIR, approaching stochastic Kriging metamodel with smoothed noise variance applied to a generic infrastructure planning problem solved by direct policy search. Our focus is on comparing the results considering continuous design variables. The problem consists of:

$$\begin{aligned} \text{Find:} \quad & \mathbf{x}^* \\ \text{that minimizes:} \quad & f(\mathbf{x}, \theta) = E[U] \\ \text{subject to:} \quad & k \in [0.05, 0.25] \\ & a_0 \in [1, 4] \\ & \Delta a \in [1, 4] \end{aligned} \tag{25}$$

Figure 2 illustrates the convergence of the expected value of the total cost, E[U], in relation to the sample size of the noisy parameter, considering the best parameter values found by Bismut and Straub (2019): $a_0 = 2$, $\Delta a = 1$, and $k = 0.1236$, resulting in a cost of 2.79. It can be observed that the expected value starts stabilizing after approximately 1600 samples. Furthermore, we observe that the expected cost, measured at the 10th and 90th percentiles, stabilizes with larger sample sizes.

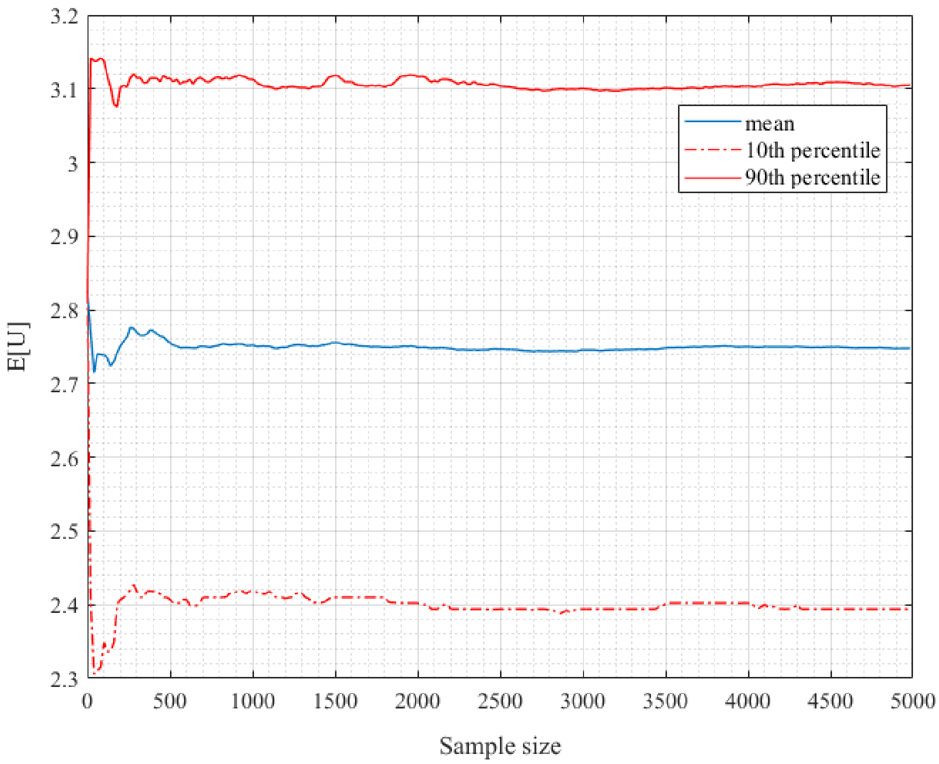

Figure 2: Expected total cost according sample size

There will be two approaches to the analysis of the problem. Initially, the direct policy search problem was simplified by considering a one-dimensional problem with only the parameter k as a design variable. Subsequently, the problem was optimized for all three parameters: a_0, Δa and k. We will evaluate both cases by varying the sampling size (n_t) of the stochastic parameter.

4.1 One-dimensional DPS problem

In this case, we considered a one-dimensional problem with k as the sole design variable. Following the results obtained by Bismut and Straub (2019), we will set the others initial parameters as $a_0 = 2$, and $\Delta a = 1$.

In the optimization process, an initial sampling $n_t = 10$ was employed to build a Kriging surrogate, distributed by Latin Hypercube method (Jones et al., 1998), while the stopping criterion was set as 3 IPs. Two sample sizes of the stochastic parameter were evaluated, $n_\theta = 50$ and $n_\theta = 100$. The performance of MQ, SEGO and EIR are presented in Table 2, with the values representing the analysis for fifty repetitions of the optimization process to account for the statistical variability.

Table 2 Optimum mean value and standard deviation for the 1D DPS problem

	$n_\theta = 50$		$n_\theta = 100$	
Optimization algorithm	Mean	Std.Dev.	Mean	Std.Dev.
MQ	2.779	0.018	2.775	0.007
sEGO	2.775	0.007	2.773	0.006
EIR	2.775	0.009	2.778	0.026

All three algorithms demonstrated strong performance, achieving average minimum values that were very close to each other. For sample sizes of $n_t = 100$, MQ, and sEGO showed lower variance in the results. The variance information approach may have influenced the results obtained for the smaller sample sizes of the stochastic parameter, $n_t = 50$, resulting in a higher standard deviation of the MQ algorithm. In contrast, the sEGO and EIR algorithms performed better, obtaining smaller minimum mean values of 2.775 and showing relatively low standard deviations. Figure 3 presents the statistical results using the box plot technique, highlighting the superior performance of sEGO and EIR for this second scenario. Furthermore, in MQ the outliers were more significant.

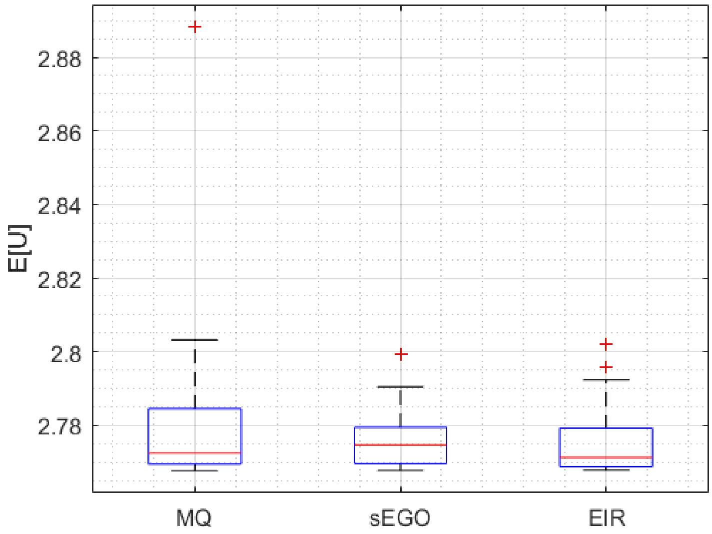

Figure 3: Boxplot of expected total cost for the 1D problem with $n_t = 50$

This analysis further highlights the strong performance of the algorithms, which achieve a mean minimum value surpassing the target of 2.79 reported by Bismut and Straub (2019) using the Cross Entropy method. Another analysis can be conducted by comparing the number of samples required to stabilize the total cost, which is 1600 as shown in Figure 2, with the number of samples used by the algorithms, which is significantly lower. The best minimum expected cost, for $n_t = 50$, reached by sEGO was 2.7677 at $k = 0.1121$and by the EIR was 2.7678 resulting from $k = 0.1125$.

4.2 Three-dimensional DPS problem

In this section we analyze the three-dimensional problem with the parameters a_0, Δa, and k as continuous design variables of heuristic DPS for the generic infrastructure planning problem.

In the process of minimizing the expected cost $E[U]$, $n_t = 30$ samples were utilized for the construction of the Kriging surrogate, along with 10 IPs in the optimization process. Two sample sizes were analyzed for the stochastic parameter, $n_\theta = 50$ and $n_\theta = 100$, with each optimization algorithm runing 50 times. The results for the mean and standard deviation are presented in Table 3.

Table 3 Optimum mean and standard deviation value for the 3D DPS problem

Optimization algorithm	Mean	Std.Dev.	Mean	Std.Dev.
MQ	2.828	0.061	2.818	0.065
sEGO	2.807	0.054	2.804	0.046
EIR	2.814	0.056	2.800	0.038

The impact of the variance information approach becomes evident in the three-dimensional DPS problem. For the two sample sizes of the stochastic parameter, the MQ algorithm resulted in a higher average of the minimums obtained. In addition, the higher standard deviation indicates greater variability in the results of the 50 iterations. The increased dimensionality of the problem may have amplified the influence of noise, affecting the algorithm's ability to handle noise heterogeneity is more important. In contrast, in a 1D approach, the influence of noise can be more easily smoothed or may have a smaller impact on the quality of the solutions.

On the other hand, the sEGO and EIR algorithms obtained a lower average value of the minimums and less variability in the results, indicating a superior performance and greater consistency in their optimization process. Among them, sEGO was the one that presented the best results, being the least affected by reducing the size of the stochastic parameter sample. These results can be evidenced when analyzing the boxplots of the optima found for $n_t = 50$, as shown in Figure 4.

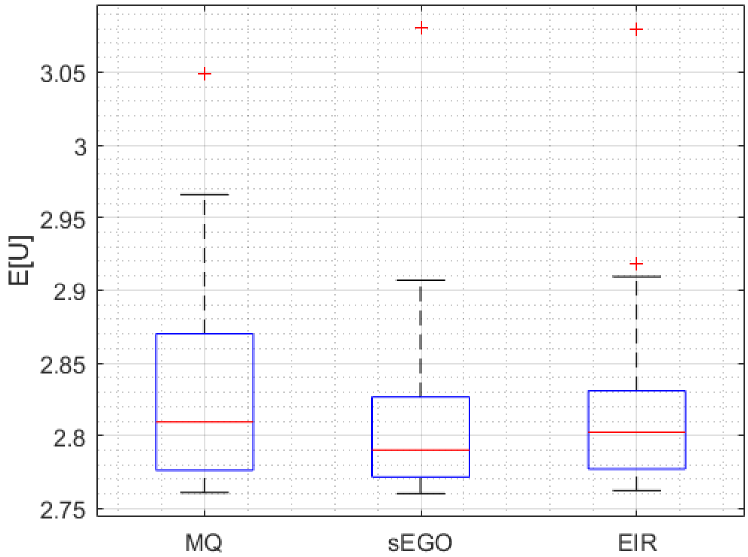

Figure 4: Boxplot of expected total cost for the 3D problem with $n_t = 50$

From Figure 4, it is possible to conclude that using the MQ algorithm resulted in greater dispersion in the results, with a higher tendency for values above the median. sEGO was the method that demonstrated the best performance, obtaining a lower average value, followed by the EIR method. Both sEGO and EIR exhibited similar dispersion in the results, with EIR having the median close to the mean value. Regarding outliers, all three methods showed a low rate of discrepancies in the minimum value. The optimal parameter values for the sEGO, and EIR algorithms, with $n_\theta = 50,$ that resulted in the minimum expected total cost, are presented in Table 4.

Table 4 Optimum parameters value for 3D DPS problem

Algorithms	E[U]	a_0	Δa	k
sEGO	2.7603	1.880	1	0.1196
EIR	2.7624	1.822	1	0.1212

5 CONCLUSION

In this paper, three algorithms - MQ, sEGO and EIR - based on Kriging for optimization were compared for solving a generic infrastructure planning problem with heterogeneous noise approaching a direct policy search heuristic. The MQ algorithm is the only one that does not address information about the noise variance of the stochastic parameter. Additionally, the study approach smoothed variance estimations in both the stochastic Kriging surrogate and the infill criterion in a sEGO and EIR optimization framework.

The problem being analyzed involves the optimization of three parameters - the initial system capacity (a_0), the capacity increment (Δa), and a capacity correction factor (k) - using a DPS heuristic as the solution strategy. Initially, the problem was simplified to a one-dimensional, with the design variable limited to only the parameter k. Later, the analysis was extended to include all three parameters as continuous design variables.

From the results, it was possible to observe that the quality of the solutions returned by the MQ algorithm was affected due to the absence of variance information in its framework. The sEGO and EIR were more competitive algorithms, presenting smaller variance in their results and achieving a lower average of the total cost obtained among all simulations from the optimization process, even with a small sample of the stochastic parameters. These results are even more evident for the three-dimensional problem, in which there was greater dispersion in the results obtained by MQ, compared to the sEGO and EIR methods.

Based on the solution obtained, it is possible to conclude that incorporating information about the variance of the stochastic parameter in Kriging-based algorithms was more effective in solving the infrastructure planning problem, using direct policy search with a heuristic approach. Furthermore, the Kriging approach not only reduced the number of expensive fitness evaluations, but also allowed optimization with a smaller sample size of the stochastic parameter without compromising the quality of the optimal solution. The results indicate that the using of Kriging-based algorithms with heterogeneous noise treatment in their framework has research potential for optimizing infrastructure planning problems using a direct policy search as a heuristic approach.

Author's Contribuitions: The authors are grateful for the support by Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) through PROAP (Graduate Support Program) - Finance Code 001 - and Federal University of Campina Grande.

Data Availability: Research data is available in the body of the article

Editor: Eduardo Alberto Fancello and Paulo de Tarso Mendonça

References

Ankenman, B, Nelson, B. L., Staum, J. (2010). Stochastic kriging for simulation metamodeling, Operations research, 58(2): 362-37.

Bismut, E., Straub, D. (2018). Adaptive direct policy search for inspection and maintenance planning in structural systems in Life Cycle Analysis and Assessment in Civil Engineering: Towards an Integrated Vision, 3103-311.

Bismut, E., Straub, D. (2019). Direct policy search as an alternative to POMDP for sequential decision problems in infrastructure planning, 13th International Conference on Applications of Statistics and Probability in Civil Engineering ICASP13, 3103-3110.

Carraro, F., Lopez, R. H., Miguel, L. F. F., Torii, A. J. (2019). Monte Carlo integration with adaptive variance selection for improved stochastic efficient global optimization, Structural and Multidisciplinary Optimization, 60, 245-268.

Cox, D. D., John, S. (1992). A statistical method for global optimization, International Conference on Systems, Man, and Cybernetics IEEE, 1241-1246.

Forrester, A. I., Keane, A. J. (2009). Recent advances in surrogate-based optimization, Progress in aerospace sciences, 45, 50-79.

Forrester, A. I., Keane, A. J., Bressloff, N. W. (2006). Design and analysis of Noisy computer experiments, AIAA, 44, 2331-2339.

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., Reed, P. M. (2016). Curses, tradeoffs, and scalable management: Advancing evolutionary multiobjective direct policy search to improve water reservoir operations, Water Resources Planning and Management, 142, 04015050.

Golla, A., Meinke, R. J., Liu, M., Staudt, P., Anderson, C. L., Weinhardt, C. (2021). Direct Policy Search for Multiobjective Optimization of the Sizing and Operation of Citizen Energy Communities. HICSS, 1-10.

Gupta, A., Liu, M., Gold, D., Reed, P., Anderson, C. L. (2020). Exploring a direct policy search framework for multiobjective optimization of a microgrid energy management system, Proceedings of the 53rd Hawaii International Conference on System Sciences.

Jalali, H., Van Nieuwenhuyse, I ., Picheny, V. (2017). Comparison of kriging-based algorithms for simulation optimization with heterogeneous noise, European Operational Research, 261, 279-301.

Jones, D. R., Schonlau, M., Welch, W. J. (1998). Efficient global optimization of expensive black-box functions, Global optimization, 13, 455-492.

Kaminski, B. (2015). A method for the updating of stochastic kriging metamodels, European Operational Research, 247, 859-866.

Luque, J ., Straub, D. (2019). Risk-based optimal inspection strategies for structural systems using dynamic Bayesian networks, Structural Safety, 76, 68-80.

Nascentes, F. F. S. (2019). Contribuições à eficiência da otimização global estocástica adaptativa, M.Sc. Dissertation (in portuguese), Federal University of Santa Catarina, Florianópolis, SC, Brazil.

Picheny, V., Ginsbourger, D., Richet, Y., Caplin, G. (2013). Quantile-based optimization of noisy computer experiments with tunable precision, Technometrics.

Quinn, J. D., Reed, P. M., Keller, K. (2017). Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points, Environmental Modelling & Software, 92, 125-141.

Rojas-Gonzalez, S., Van Nieuwenhuyse, I. (2020). A survey on kriging-based infill algorithms for multiobjective simulation optimization, Computers & Operations Research, 116, 104869.

Van Beers, W. C., Kleijnen, J. P. (2003). Kriging for interpolation in random simulation, Operational Research Society, 54, 255-262.