Key Engineering Materials
ISSN: 1662-9795, Vol. 1049, pp 141-150
doi:10.4028/p-1kSPJP

Submitted: 2026-02-02
Revised: 2026-02-24
Accepted: 2026-03-02
Online: 2026-04-15

Robust Performance Optimization Using Bayesian Optimization and Extreme Value Theory

Augustin Persoons[a,*], Valentin Duarte Rocha[b], Laurence Giraud Moreau[c], Pascal Lafon[d]
Université de Technologie de Troyes, Unité de Recherche LASMIS
[a]augustin.persoons@utt.fr, [b]valentin.duarte_rocha@utt.fr, [c]laurence.moreau@utt.fr, [d]pascal.lafon@utt.fr

Keywords: Bayesian Optimization, Optimization Under Uncertainty, Robustness, Active Learning, Forming Processes.

Abstract. This work deals with Robust design optimization (RDO) under interval uncertainty and the resolution of such problems using a Bayesian optimization algorithm. In metal forming, process parameters such as tool radius, step size, or forming toolpath introduce variability that directly affects the final geometry and quality of the formed parts. In this context we aim at finding a design minimizing the amplitude of the performance interval but such a formulation does not account for the nominal performance. In this work, we introduce a scalarized objective adapted to the proposed algorithm allowing it to identify a Pareto optimum of both stability and nominal behavior. We propose an efficient expected improvement (EI) estimator for this objective based on an extreme-value approximation of surrogate extrema. The approach is illustrated on an analytical test problem and on a forming simulation with spring-back, where the new objective yields more practically relevant solutions than a variation-only robustness criterion.

1 Introduction and Problem formulation

In many industrial settings, the performances of a system or process must not only perform well at nominal conditions, but must also remain acceptable under variability of operating conditions, material properties, or manufacturing parameters. This motivates robust design optimization (RDO), i.e. the search for design variables that reduce sensitivity to uncertainty while meeting performance requirements. Robustness is often associated with Taguchi's work [1], and a broad spectrum of formulations has been proposed since then, ranging from stochastic approaches based on moments or failure probabilities to deterministic worst-case approaches [2–4].

This work considers the case where uncertain variables and design variables are distinct. Moreover, uncertainties are described by intervals rather than probability distributions. Such a setting occurs when uncertainty data is scarce, heterogeneous, or not sufficiently stationary to justify a probabilistic model, yet reliable bounds can be specified based on expert's knowledge. Let $x \in I_x \subset \mathbb{R}^{n_x}$ denote design variables for which we want to optimize the system and $u \in I_u \subset \mathbb{R}^{n_u}$ uncertain variables which can take any value within the interval, and let the (expensive) performance model be

$$y = g(x, u)$$

Since the parameter set u is an interval, the performance for any design x is also an interval. A classical robustness metric over interval uncertainty is the range (worst-case spread) of the response,

$$\Delta_y(x) = \max_{u \in I_u} g(x, u) - \min_{u \in I_u} g(x, u)$$

which quantifies the amplitude of performance variations associated with a given design (x). A first step toward robust design optimization could be to focus on minimizing $\Delta_y(x)$, to reach the design with the most stable performance. This is, however, a two-level nested optimization problem which is notoriously computationally demanding and usually untractable if g is computationally expensive to evaluate. For the sake of clarity let us refer to the problem of identifying, for a given design x, the most extreme performance values $\max_{u\in I_u} g(x,u)$ and $\min_{u\in I_u} g(x,u)$ as the inner-problem. Likewise, let us refer to the problem of identifying which design is the most stable ($\min_{x\in I_x} \Delta_y(x)$) as the outer-problem.

Previous work focused on proposing to solve this RDO problem with Bayesian optimization first based on confidence interval [5] and then using a heuristic based on extreme-value distributions [6]. However, optimizing $\Delta_y(x)$ alone may yield designs with stable but consistently poor performance. For this reason, the present paper revisits the objective formulation and proposes a scalar mean-variation criterion that can be interpreted as a compact alternative to an explicit two-objective formulation.

Let us define, for any design x, the extreme responses

$$\bar{g}(x) = \max_{u\in I_u} g(x,u), \qquad \underline{g}(x) = \min_{u\in I_u} g(x,u),$$

and the mid-range (center of the response interval)

$$\mathrm{m}_y(x) = \frac{\bar{g}(x) + \underline{g}(x)}{2}.$$

We define the objective

$$J(x) = \frac{1}{2}\sqrt{\Delta_y(x)^2 + \left(2\,\mathrm{m}_y(x)\right)^2} \tag{1}$$

assuming the design optimization problem is formulated such that the optimal performance is associated with $g(x,u) = 0$ (e.g. offset from a desired performance value), this criterion penalizes both spread and offset in a single quantity. Using $\Delta_y(x) = \bar{g}(x) - \underline{g}(x)$, it follows that minimizing $J(x)$ is equivalent to minimizing the smooth quadratic form

$$J^2(x) = \bar{g}(x)^2 + \underline{g}(x)^2 \tag{2}$$

up to a monotone transformation.

The problem is illustrated in Figure 1 on the analytical benchmark 2D function described in section section 4. Figure 5a illustrates the performance surface as a function of design and uncertain parameters x, u. Figure 1b is a projection of that surface on the $x - y$ plane, illustrating the amplitude of variation of the performance over u and how this amplitude evolves as a function x. Figure 1c represents the evolution of Δ_y as a function x, also highlighting the optimal design with respect to this criterion. Figure 1c represents the evolution of m_y as a function x and associate optimum. Figure 1d represents the evolution of the multi-objective criterion J as a function x with associated optimal design.

The computational challenge is that each evaluation of $J(x)$ (or $J^2(x)$) involves two inner optimization problems over u, and the outer optimization over x must be performed with a limited simulation budget. Bayesian optimization (BO) addresses such settings by replacing the expensive model g with a surrogate model $\hat{g}$ (a Gaussian process regression model, GPR). The surrogate is calibrated with an active learning strategy aiming at minimizing the number of calls to g necessary to estimate the optimum. The active calibration scheme uses a so-called acquisition function (e.g. expected improvement, EI) to identify which sample to include next in the calibration dataset [7].

The main difficulty in this context is to propagate the stochasticity of the surrogate $\hat{g}$ to the objective Δ_y to guide the active calibration. In standard BO, EI is straightforward to compute when the objective at a candidate point is Gaussian. In the present setting, the objective depends on extrema of the Gaussian process, and accurately characterizing the distribution of these extrema is known to be challenging for non-stationary Gaussian processes [8, 9]. A naive solution based on Monte Carlo sampling of GPR trajectories would be prohibitively expensive when embedded in an outer optimization

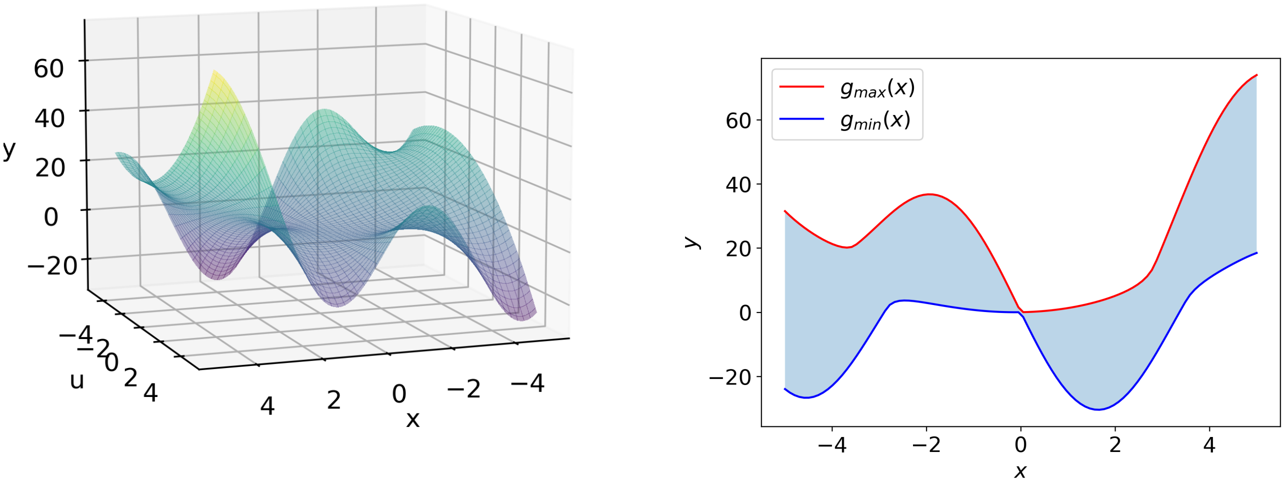

(a) Example performance surface as a function of design and uncertain parameters

(b) Projection of the performance surface on the design axis

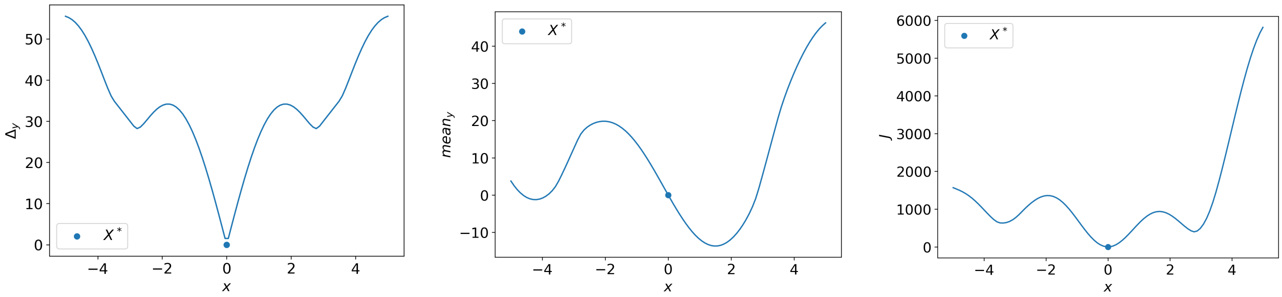

(c) Performance amplitude as a function of x

(d) Performance average as a function of x

(e) Multi-objective criterion as a function of x

Figure 1: Illustration of the RDO problem and associated quantities of interest

loop. The proposed approach relies on building a simplified stochastic representation of Δ_y using extreme value distributions in two steps. First, we approximate the distributions of the extrema induced by a virtual population of GPR realizations using extreme-value theory (Gumbel-type approximations). Second, exploiting the quadratic form in eq. (2), using a simple change of variables we can numerically estimate the EI through numerical integration at a moderate and controllable cost. The resulting method retains the practical benefits of EI-driven exploration–exploitation while avoiding repeated sampling of full GPR trajectories.

The rest of the paper is organized as follows: Section 2 introduces the GPR surrogate setting and the extreme-value approximation of its extrema. Section 3 develops the acquisition function proposed and the Bayesian optimization scheme Finally, Section 4 presents the numerical examples including the forming simulation, and Section 5 concludes with perspectives for future work.

2 Gaussian-Process Surrogate and Extreme-Value Approximations of Its Extrema

2.1 Gaussian-process regression model The goal of this section is to build a surrogate of the expensive model $g(x, u)$, using a Gaussian process regression model (GPR). Their principle is to consider the true performance surface as a realization of a Gaussian process (GP). By assuming a GP prior, a covariance function, and given some data points on the true performance surface, the prior can be conditioned to the available data to calibrate a predictive posterior through Bayesian inference.

For a more details on GPRs, interested readers can refer to [10]. In practice this type of surrogate can be trained in a similar fashion as any regression model, the main difference is that their prediction is stochastic in nature. Meaning that, for any candidate input $(x, u) \in I_x * I_u$, the predicted performance is a normally-distributed random variable with mean $\mu(x, u)$ and standard deviation $\sigma(x, u)$,

$$\hat{g}(x, u) \sim \mathcal{N}\big(\mu(x, u),\, \sigma(x, u)\big) \tag{3}$$

The mean value can be seen as the most likely performance values according to the surrogate, while the standard deviation can be seen as the surrogate's prediction uncertainty.

The outer robust objective introduced in section 1 depends on the true (model-based) extrema

$$\bar{g}(x) = \max_{u \in I_u} g(x, u), \qquad \underline{g}(x) = \min_{u \in I_u} g(x, u),$$

and the corresponding quantity to be minimized is

$$J(x) = \bar{g}(x)^2 + \underline{g}(x)^2$$

As previously mentioned, these quantities are not available to us and are too computationally expensive to evaluate directly. However, once a GP posterior $\hat{g}$ is available, on can define the surrogate extrema defined as

$$\bar{\hat{g}}(x) = \max_{u \in I_u} \hat{g}(x, u), \qquad \underline{\hat{g}}_N(x) = \min_{u \in I_u} \hat{g}_N(x, u). \tag{4}$$

A key point is that since $\hat{g}(\cdot)$ is a Gaussian process, $\bar{\hat{g}}_N(x)$ and $\underline{\hat{g}}_N(x)$ are random variables for every fixed x. Consequently, the surrogate-based objective

$$\hat{J}(x) = \bar{\hat{g}}(x)^2 + \underline{\hat{g}}(x)^2 \tag{5}$$

is itself a random variable, meaning that $\hat{J}$ is a stochastic process and similarly to $\hat{g}$, its variability is related to the surrogate's prediction uncertainty.

In other words, one can see $\hat{J}$ as a stochastic surrogate of the RDO objective function. The idea is then to use a Bayesian optimization strategy applied to this surrogate and thus solve the RDO problem.

2.2 The expected improvement acquisition function One of the most popular Bayesian optimization method is the efficient global optimization methods (EGO) [7]. Which proposed a heuristic for selecting the coordinates to refine named the expected improvement.

Let us consider an arbitrary step N of the Bayesian optimization algorithm and the current best robust value available to us $J^\star_N$. Note that this value is a deterministic scalar, not a random variable. Given the current calibration of the GPR surrogate, let us define the improvement of a design as the difference (bounded by zero) between the objective values J of the current optimum and the candidate design x.

$$I_N(x) = \max\Big(0,\ J^\star_N - \hat{J}(x)\Big),$$

since $\hat{J}(x)$ is a random variable, $I_N(x)$ is also one. The corresponding expected improvement (EI) criterion is then defined as:

$$\mathrm{EI}(x) = \mathbb{E}[I(x)] = \mathbb{E}\Big[\max\Big(0,\ J^\star - \big(\bar{\hat{g}}(x)^2 + \underline{\hat{g}}(x)^2\big)\Big)\Big]\,. \tag{6}$$

In a standard setting (where the objective is directly the GP prediction), EI_N admits a closed form. The issue in the RDO case is that the distributions of $\bar{\hat{g}}$ and $\underline{\hat{g}}$ are not available to us. They are random variables defined as the extrema of a Gaussian process so there is not closed form for their distributions in the general case and eq. (6) cannot be evaluated analytically. Therefore we need a practical and computationally tractable approximation strategy for the distributions of $\bar{\hat{g}}$ and $\underline{\hat{g}}$.

2.3 A heuristic based on extreme value distributions This section summarizes the key idea used to obtain a tractable approximate distribution for the surrogate-based extrema $\bar{\hat{g}}$ and $\underline{\hat{g}}$. Which are needed to compute acquisition functions such as Expected Improvement (EI) for robust objectives. A more detailed derivation and empirical justification can be found in [6].

The heuristic relies on the following idea: for a fixed (x, u), consider a virtual population of n independent draws from the predictive distribution $Y = \hat{g}(x, u)$

$$Y_1, \dots, Y_n \overset{iid}{\sim} \mathcal{N}\big(\mu(x, u),\, s^2(x, u)\big),$$

and denote their maximum by

$$Y^{(n)}_{\max}(x, u) = \max_{1 \le i \le n} Y_i.$$

This extrema is a random variable and extreme-value theory suggests that, for moderately large n, the law of $Y^{(n)}_{\max}$ is well approximated by a Gumbel distribution after affine normalization. In practice, we only need a fast approximation of the Gumbel location and scale parameters that can be directly extracted from the surrogate posterior values $\mu(x, u)$ and $s^2(x, u)$ and depend only on n. A distribution of the minimum Y^n_{min} can be similarly approximated.

For any given design x, we then use a statistics (e.g. mean or a quantile) of this extreme value distribution as a heuristic score to locate where the extremum over u is most likely to occur (locate the maximizer and minimizer):

$$u^+(x) \simeq \operatorname*{argmax}_{u \in I_u} \Big(E(Y^n_{max})\Big), \qquad u^-(x) \simeq \operatorname*{argmin}_{u \in I_u} \Big(E(Y^n_{min})\Big). \tag{7}$$

Intuitively, eq. (7) balances exploitation (large/small mean) and exploration (large uncertainty), and returns two representative locations: one for the maximum and one for the minimum.

Once the likely maximizer and minimizer u values $(u^+(x), u^-(x))$ are located, we approximate the extrema of the Gaussian process $\hat{g}(x, u)$ over u by the virtual extremes at these locations:

$$\bar{\hat{g}}_N(x) \approx Y^{(n)}_{\max}\big(x, u^+(x)\big), \qquad \underline{\hat{g}}_N(x) \approx -Y^{(n)}_{\max}\big(x, u^-(x)\big). \tag{8}$$

This yields two explicit one-dimensional approximate distributions, which are the only objects required by the next-step EI computation. It is however worth noting that this represents a simplistic approximation of the actual distribution only meant to guide the calibration of the surrogate in a pseudo-optimal manner. This simplification means in essence reducing the Gaussian process to a single Gaussian random variable, the one corresponding to the maximizer/minimizer location and cannot be considered as a proper approximation of the extrema of the GP.

3 Acquisition Function and Bayesian Optimization Algorithm

It is important to distinguish the three levels of approximation of the objective function:

- the true robust objective

  $$J(x) = \bar{g}(x)^2 + \underline{g}(x)^2, \qquad \bar{g}(x) = \max_{u \in I_u} g(x, u), \quad \underline{g}(x) = \min_{u \in I_u} g(x, u),$$

  which can only be accessed through expensive calls to g

- the surrogate-induced stochastic objective

  $$\hat{J}(x) = \bar{\hat{g}}(x)^2 + \underline{\hat{g}}(x)^2,$$

  where $\bar{\hat{g}}(x)$ and $\underline{\hat{g}}(x)$ are random because $\hat{g}$ is a posterior GP. In the proposed method, calibration decisions (where to sample next) are made by maximizing an acquisition function requiring the estimation of statistics of $\hat{J}$. Unfortunately, as mentioned in the previous section, the distribution of $\hat{J}$ doesn't have a closed-form expression.

- the extreme-value heuristic used as a computationally tractable proxy for $\hat{J}$

$$\hat{J}_{ev}(x) = Y^{(n)}_{\max}\big(x, u^{+}(x)\big)^2 + Y^{(n)}_{\min}\big(x, u^{-}(x)\big)^2,$$

as described in equations 8 and 7.

3.1 Expected improvement under the extreme-value approximation To obtain a tractable EI, we assume Y^n_{max} and Y^n_{min} are independent under the fitted models. This is an extra approximation: maxima and minima of the same random field are generally dependent. Empirically, the acquisition remains effective because EI is used only as a ranking criterion over x. Then, considering a current optimal value $J^\star$ the expected improvement is approximated as:

$$\mathrm{EI}_N(x) \approx \mathbb{E}\Big[\max\Big(0,\ J^\star - Y^{(n)}_{\max}\big(x, u^{+}(x)\big)^2 + Y^{(n)}_{\min}\big(x, u^{-}(x)\big)^2\big)\Big)\Big]. \tag{9}$$

In the multi-objective robust criterion used here, let us consider the radius $R = \sqrt{(Y^n_{max})^2 + (Y^n_{min})^2}$ and define the improvement w.r.t. the current best estimate $J^\star$ by $I(J^\star) = (J^\star - R)$. The Expected Improvement is then

$$\mathrm{EI}(\tau) = \mathbb{E}\big[I(J^\star)\big] = \int_0^{J^\star} (J^\star - r)\, f_R(r)\, dr, \tag{10}$$

where f_R is the density of R. Under independence, the joint density is $f_{X,Y}(x,y) = f_X(x) f_Y(y)$, and by the polar change of variables $(x,y) = (r\cos\theta, r\sin\theta)$ (Jacobian r) we obtain

$$f_R(r) = r\int_0^{2\pi} f_X(r\cos\theta)\, f_Y(r\sin\theta)\, d\theta, \qquad r \geq 0.$$

We evaluate the inner angular integral numerically and then compute $\mathrm{EI}(J^\star)$ by quadrature on $[0, J^\star]$.

3.2 Overall procedure The next section reports numerical experiments and compares the proposed strategy against baselines (e.g. standard EI on a scalarized objective, worst-case UCB, and random sampling), using a fixed evaluation budget.

Since the algorithm estimates the amplitude of performance y for any given design x, at the next candidate sample for refinement x_{N+1} there is two candidate u positions, u^{+}_{N+1} and u^{-}_{N+1} as described in section 2.3. In order to limit the number of calls to the expensive model g, only one of the two candidate points is actually evaluated and used to enrich the calibration dataset. The choice is made based on the variables of the Gumbel distributed variables Y^n_{max} and Y^n_{max} where the sample being associated with the highest variance is selected.

The overall method procedure is finally described in algorithm 1.

4 Examples and Results

4.1 Analytical benchmark 1: local minima The convergence of the proposed method is first assessed on a simple two-dimensional analytical benchmark chosen to exhibit several local minima as illustrated in figure 1. The test case is defined by the following performance function g:

$$g(x,u) = xu - \sin(x)\, u^2 + x^2, \qquad (x,u) \in I_x \times I_u. \tag{11}$$

The design variable x is restricted to the interval $I_x = [-5, 5]$, while the uncertain variable u varies within $I_u = [-5, 5]$. For this benchmark, the reference optimal design is $x^* = 0$.

The convergence of the proposed method on this example is shown in Figure 2. The figure illustrates that the algorithm was able to stably reach the global minimum after about 20 iterations As a perspective, a related method previously developed by the authors required around 27 iterations to reach convergence on this example [5]. While requiring further statistical analysis, this resultis a first indicator of the good performances of this proposed approach.

Algorithm 1 Bayesian optimization for interval-robust design via extreme-value EI

```
Choose an initial design set $\{(x_i, u_i)\}_{i=1}^{N_0}$ and evaluate $y_i = g(x_i, u_i)$.
Fit a GP surrogate $\hat{g}_{N_0}$ on $\mathcal{D}_{N_0}$.
Initialize $J^\star$ from robust evaluations of a small set of designs.
for $N = N_0, N_0 + 1, \dots$ until budget is exhausted do
    For candidate $x$, fit Gumbel models for $\bar{\hat{g}}_N(x)$ and $\underline{\hat{g}}_N(x)$ (Section 2.3).
    Compute $\mathrm{EI}_N(x)$ via eq. (10).
    Select $x_{N+1} \in \mathrm{argmax}_{x \in I_x} \mathrm{EI}_N(x)$.
    Choose $u^+_{N+1}, u^-_{N+1}$ as described in section section 3.1
    Evaluate $g(x_{N+1}, u^+_{N+1})$ or $g(x_{N+1}, u^-_{N+1})$; update $\mathcal{D}_{N+1}$.
    Refit the GP to obtain $\hat{g}_{N+1}$.
    Estimate $\hat{J}^\star(x_{N+1})$ by direct estimation using $Y_{max}$, $Y_{min}$ and update $J^\star_{N+1}$.
end for
return Best design $x^\star$ associated with $J^\star_{2,N}$.
```

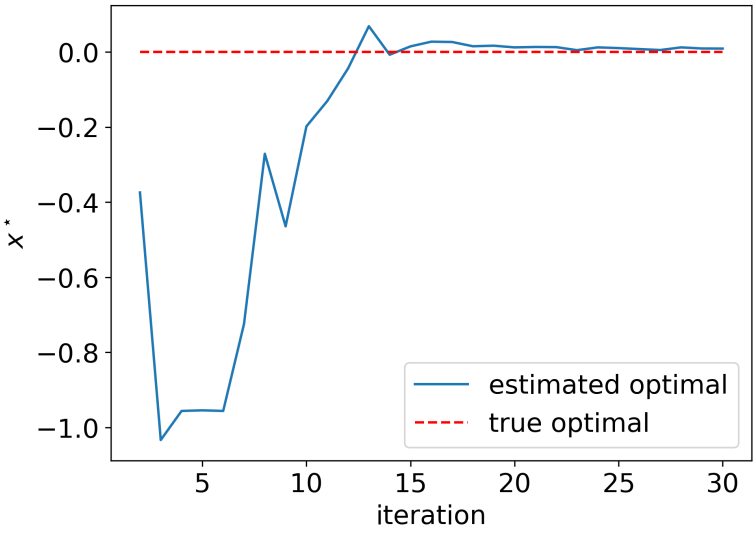

Figure 2: Convergence of the algorithm on the analytical benchmark case

4.2 Forming process finite element example with springback The last example is based on a 2D finite element simulation of a sheet-metal forming process, in which a metal blank is clamped against the die by a holder and bent into a "U" shape by a punch (see [11]). The material is DP780 steel with a thickness of $1.4\,\mathrm{mm}$. After forming, the tool is released and the elastic recovery is simulated in order to account for springback. The quantity of interest is the final sheet profile after springback.

The profile is assumed to depend on two parameters: the die-edge rounding radius R and the force applied to the Holder F. In this study, F is treated as an interval uncertainty with $I_F = [1, 10] \times 10^3\,\mathrm{N}$ while R is the design variable with design space $I_R = [1, 10]\,\mathrm{mm}$. In essence the problem can be seen as identifying which die geometry is the most robust to holding force variation. This example should be seen as a simple benchmark example and is not intended to represent realistic industrial challenges.

The performance metric is defined as the root mean square error (RMSE) between the nodal coordinates of the sheet profiles before and after springback (see an illustration in Figure 4).

The problem can be illustrated in a similar fashion as the first one in Figure 3. The contrary to the actual application of the method, the illustration is based on GPR surrogate of the finite element model, built with an previously performed design of experiment (red dot in figure). This result is only used for illustration purposes and to obtain easily computable insights of the the model behavior.

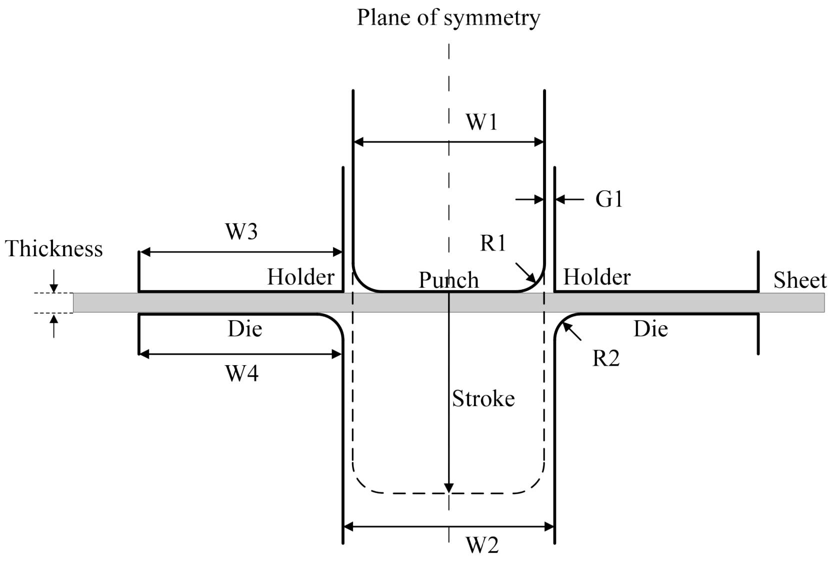

Parameters	W1	W2	W3	W4	R1	R2	G1	Stroke
Dimensions	50.0	54.0	89.0	89.0	[2, 5]		2.0	71.8

(Unit: mm)

Figure 3: Schematic representation and dimensions of the forming process.

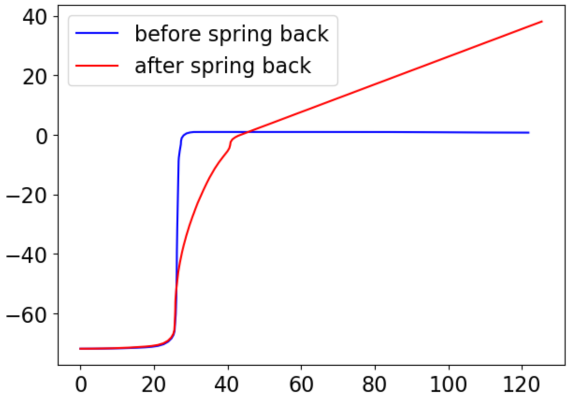

Figure 4: Illustration of the formed sheet profile before and after springback.

As illustrated in Figure 3, the performance surface of the finite element model is relatively well behaved without being trivial. The problem seem to exhibit clear optimal designs $x^\star$ in terms of amplitude, nominal value and multi objective indicator.

The proposed method was run for 20 iterations on the finite-element model and the convergence is shown in Figure 6. The method was initialized with a latin-hypercube design of experiment of 10 evaluations og the finite-element model. The method evidently converged the expected optimal in a few iterations which is a good first indicator of the performances of the method.

5 Conclusions and Perspectives

This paper proposes a Bayesian optimization strategy for robust design optimization under interval uncertainty, targeting expensive black-box simulations. In contrast with our previous work, which emphasized a variation-only robustness measure, we introduce a mean–variation objective that explicitly balances nominal performance and robustness.

To enable efficient optimization with a Gaussian process surrogate, we develop an expected improvement (EI) estimation approach tailored to this extrema-based objective. The method combines (i) an extreme-value modeling of the maxima and minima of GP trajectories and (ii) a numerical integration scheme made tractable by a polar-like change of variables. This construction yields a

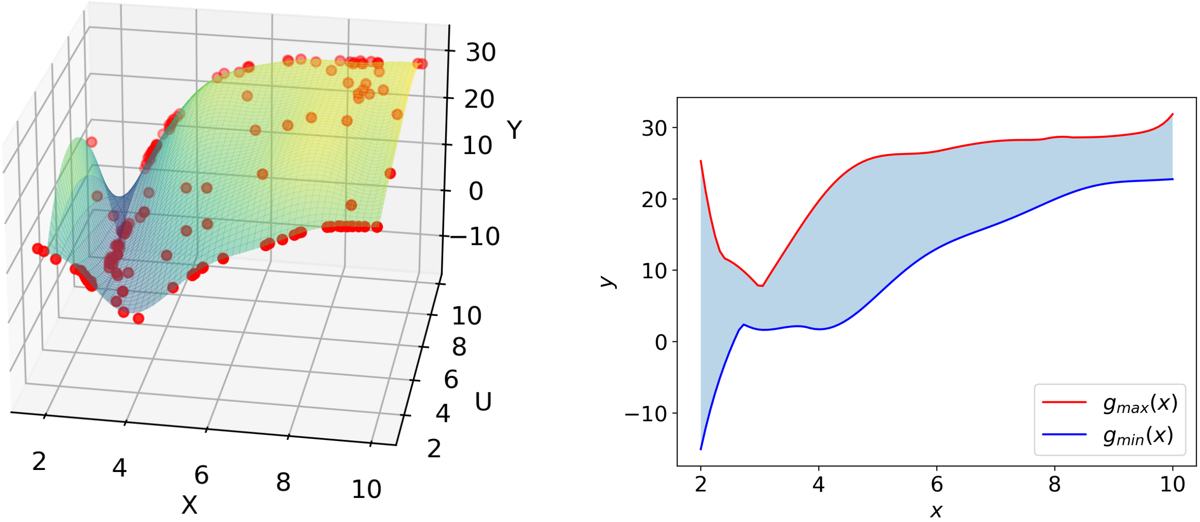

(a) Performance surface of the FE model as a function of design and uncertain parameters and points used to calibrate the response surface

(b) Projection of the FE performance surface on the design axis

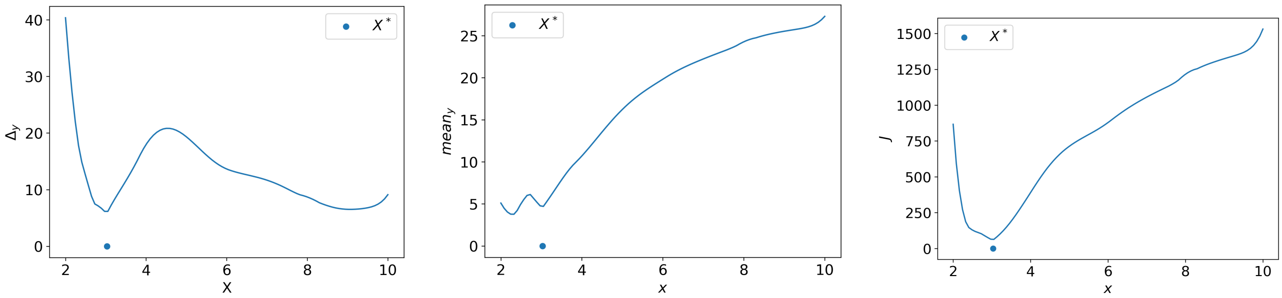

(c) Performance amplitude of the FE model as a function of x

(d) Performance average of the FE model as a function of x

(e) Multi-objective criterion of the FE model as a function of x

Figure 5: Illustration of the RDO problem and associated quantities of interest

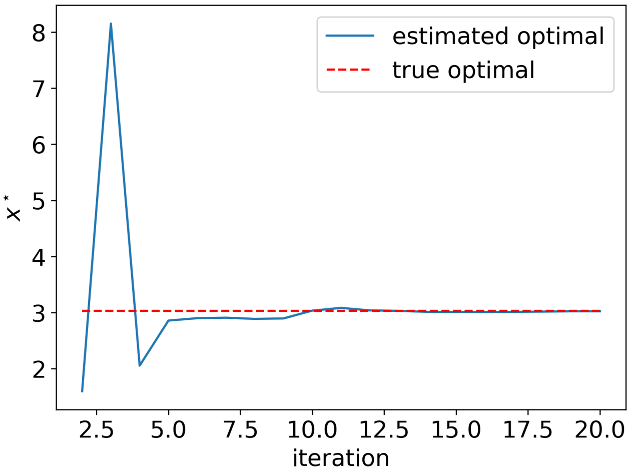

Figure 6: Convergence of the proposed algorithm on the finite element problem.

computationally efficient EI approximation that remains closely connected to the robustness metric of interest.

The approach is illustrated on two examples: an analytical test problem featuring local minima and a finite element forming simulation with springback. In both cases, the method is able to drive the search toward robust designs while accounting for nominal performance, and it provides solutions that are more practically relevant than those obtained with a variation-only robustness criterion.

An application of this method on optimizing an incremental forming process is being investigated. Future work will focus on a systematic sensitivity analysis of key modeling and numerical choices (notably the extreme-value approximation and the integration settings), as well as extended comparisons on a broader and fully reproducible benchmark suite and additional industrial-scale simulation cases.

References

[1] G. Taguchi, Quality engineering (Taguchi methods) for the development of electronic circuit technology, *IEEE Transactions on Reliability*. 44 (June 1995) 225–229.

[2] H.-G. Beyer, B. Sendhoff, Robust optimization – A comprehensive survey, *Computer Methods in Applied Mechanics and Engineering*. 196 (July 2007) 3190–3218.

[3] C. Zang, M. Friswell, J. Mottershead, A review of robust optimal design and its application in dynamics, *Computers & Structures*. 83 (Jan. 2005) 315–326.

[4] D. Bertsimas, D. B. Brown, C. Caramanis, Theory and Applications of Robust Optimization, *SIAM Review*. 53 (Jan. 2011) 464–501.

[5] C. van Mierlo, A. Persoons, M. G. R. Faes, D. Moens, Robust Design Optimization of Expensive Stochastic Simulators Under Lack-of-Knowledge, *ASCE-ASME J Risk and Uncert in Engrg Sys Part B Mech Engrg*. 9 (June 2023) 021205.

[6] A. Persoons, P. Lafon, D. Moens, BAYESIAN ROBUST DESIGN OPTIMIZATION: AN EXTREME VALUE APPROACH, *Uncecomp 2025 proceedings*. (June 2025).

[7] D. R. Jones, M. Schonlau, Efficient Global Optimization of Expensive Black-Box Functions, *Journal of Global Optimization*. 13 (1998) 455–492.

[8] J.-M. Azaïs, L.-V. Lozada-Chang, A toolbox on the distribution of the maximum of Gaussian process, (2013).

[9] G. Lindgren, Gaussian Integrals and Rice Series in Crossing Distributions—to Compute the Distribution of Maxima and Other Features of Gaussian Processes, *Statistical Science*. 34 (Feb. 2019).

[10] R. B. Gramacy, Surrogates: Gaussian Process Modeling, Design and Optimization for the Applied Sciences, Chapman Hall/CRC, Boca Raton, Florida, 2020.

[11] S. H. e. W. C. GH. Huh K. Chung, Benchmark Study of the 8th International Conference and Workshop on Numerical Simulation of 3D Sheet Metal Forming Processes, *Proceedings of Numisheet 2011*. (2011).