Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Configuration-Space Bayesian Optimisation for Stress-Critical Pose Identification on a Hyper-Redundant Robot

Ardit Poka

ardit.poka@iit.it

Italian Institute of Technology

Daniele Ludovico

Italian Institute of Technology

Lorenzo De Mari Casareto Dal Verme

Italian Institute of Technology

Federico Manara

Italian Institute of Technology

Carlo Canali

Italian Institute of Technology

Giovanni Berselli

University of Genoa

Darwin G. Caldwell

Italian Institute of Technology

Research Article

Keywords: Bayesian optimisation, Robotics, Surrogate modelling, Hyper-redundant robots, Finite Element Analysis

Posted Date: June 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9542708/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at International Journal of Intelligent Robotics and Applications on September 3rd, 2026. See the published version at https://doi.org/10.1007/s41315-026-00586-z.

Configuration-Space Bayesian Optimisation for Stress-Critical Pose Identification on a Hyper-Redundant Robot

Ardit Poka · Daniele Ludovico · Lorenzo De Mari Casareto Dal Verme · Federico Manara · Carlo Canali · Giovanni Berselli · Darwin G. Caldwell

Abstract Serial manipulators can have strongly varying internal stresses within their configuration space, making the identification of stress-critical poses important for structural design. Because each candidate configuration requires a computationally expensive Finite Element Analysis (FEA), high-fidelity simulations quickly become impractical. This paper proposes a Bayesian optimisation framework that searches directly over the robot configuration space, using peak Von Mises stress as the objective. For each candidate configuration, statically consistent interface reactions are computed and applied as boundary conditions in the structural model. A Gaussian process surrogate with an ARD Matérn $5/2$ kernel models the configuration-stress relationship, while the Bayesian algorithm with Expected Improvement acquisition function selects new configurations for parallel evaluation. A case study on a cable-driven hyper-redundant robot demonstrates that the proposed workflow can identify configurations with critical stress states and accurately map the high-stress region with a limited number of simulations. The results show that the framework can support both structural verification and the detection of potentially critical robot poses during operation.

Corresponding author: Ardit Poka

Ardit Poka, Daniele Ludovico, Federico Manara, Lorenzo De Mari Casareto Dal Verme, Giovanni Berselli, Carlo Canali, and Darwin G. Caldwell
Advanced Robotics (ADVR), Italian Institute of Technology

Giovanni Berselli
Dept. of Mechanical, Energy, Management and Transportation Engineering (DIME), University of Genova

E-mail:
ardit.poka@iit.it, daniele.ludovico@iit.it,
lorenzo.demaricasareto@iit.it federico.manara@iit.it
carlo.canali@iit.it, giovanni.berselli@unige.it, darwin.caldwell@iit.it,

Keywords Bayesian optimisation · Robotics · Surrogate modelling · Hyper-redundant robots · Finite Element Analysis

1 Introduction

Structural design optimisation aims to use minimal material while meeting performance constraints, resulting in lightweight, strong structures [1]. Common approaches include topology optimisation to distribute material within a domain, and shape/size optimisation for components. Evolutionary and heuristic algorithms are widely used for high-dimensional or nonconvex design spaces, but they often require many function evaluations [2]. More recently, machine-learning surrogate models and hybrid methods have been introduced to predict performance and to accelerate optimisation by reducing the number of simulation runs. These advances have allowed practical optimisation in engineering domains such as aerospace, automotive, and civil structures, balancing objectives such as weight minimisation, strength, and sustainability [1, 3].

In robotic and multibody system design, optimisation must account for coupled dynamics, kinematics, and control. For example, Tromme et al. [4] couple a dynamic multibody simulation with gradient-based level-set optimisation of individual components, so the component shapes are optimised under realistic system loads and motions. As robot morphology becomes more articulated, this system-level coupling becomes even more pronounced: small changes in posture can substantially alter contact conditions, cable routing, and internal load paths. Hyper-Redundant (HR) manipulators, characterised by a high number of Degrees of Freedom (DOF) [5], are a representative and increasingly important class of such systems, offering a promising solution for tasks requiring manoeuvrability and adaptability in constrained environments [6–8]. However, the high-DOF configuration

space induces a very large set of possible structural states, and different feasible robot poses can produce significantly different internal stresses. This makes the identification of stress-critical configurations important for structurally robust design. Identifying these critical configurations is challenging because the mapping from robot configuration to peak stress is highly nonlinear and typically accessible only through computationally intensive simulations. Therefore, efficient optimisation strategies are needed to explore this large space and locate worst-case responses with as few evaluations as possible.

To address this need, this paper presents a method for worst-case stress identification for a cable-driven hyper-redundant manipulator as a black-box maximisation problem over the feasible robot configuration space, where each function evaluation requires a high-fidelity FEA simulation. For each candidate configuration, the robot statics model is solved to compute the corresponding interface reactions, which are then applied as boundary conditions in the structural analysis of the component of interest. Bayesian optimisation with a Gaussian process surrogate model is employed to adaptively select informative configurations, concentrating samples in high-stress regions while maintaining exploration, thereby reducing the number of expensive simulations required to locate the worst-case feasible configuration.

Bayesian Optimisation (BO) is a gradient-free, probabilistic framework for optimising expensive black-box functions [9]. BO builds a surrogate model, typically a Gaussian Process (GP), of the unknown objective and iteratively refines it by sampling promising designs [10]. Because GPs quantify uncertainty, BO can selectively propose samples that either improve the best-known solution or reduce uncertainty around high-variance regions. This makes BO especially useful when each evaluation (e.g. FEA) is costly or noisy. In recent years, BO and GP surrogates have been successfully applied to structural and robotic design. For instance, Tran et al. [11] used BO to drive a moving-morphable-bar topology optimisation, finding better solutions than gradient-based methods in highly nonlinear design spaces while requiring fewer analyses. BO has also been used for co-design of mechanics and behavior [12]. BO/GP methods are particularly attractive for expensive black-box engineering problems because they combine surrogate modelling with uncertainty-aware sampling. In settings where each function evaluation is costly, such as FEA-based structural analysis, this makes them a practical choice for guiding the search toward informative regions of the design or operating space [13, 14]. In addition, Gaussian processes are well-suited to low-data regimes and provide predictive uncertainty, which is useful when the objective is not only to approximate the response surface, but also to identify potentially critical regions that may otherwise remain unexplored [9, 10]. For these reasons, this work adopts a BO-GP framework as a practical and physically interpretable tool for worst-case stress identification in robot configuration space.

Table 1: DH parameters of the robot.

Joint	$\boldsymbol{a}$ [mm]	$\boldsymbol{\alpha}$ [rad]	$\boldsymbol{d}$ [mm]	$\boldsymbol{\theta}$ [deg]
1	0	$\pi/2$	0	$q_{1,1}$
1_c	100	0	0	$q_{2,1}$
2	100	$-\pi/2$	0	$q_{1,2}$
2_c	100	0	0	$q_{2,2}$
3	100	$\pi/2$	0	$q_{1,3}$
3_c	100	0	0	$q_{2,3}$
4	100	$-\pi/2$	0	$q_{1,4}$
4_c	100	0	0	$q_{2,4}$
5	100	$\pi/2$	0	$q_{1,5}$
5_c	100	0	0	$q_{2,5}$
EE	100	0	0	0

The goal of this work is not to provide a broad benchmark of optimisation algorithms or surrogate modelling techniques. Rather, the aim is to demonstrate that a BO-GP workflow can serve as a practical tool for identifying stress-critical feasible configurations in a cable-driven hyper-redundant robot, where each structural evaluation is computationally expensive. Accordingly, the validation focuses on physical consistency, surrogate accuracy, convergence toward critical configurations, and the framework's ability to support worst-case structural assessment within a constrained FEA simulation. It is important to note that the purpose of the present case study is methodological, and the first link considered here is intentionally overdesigned and therefore exhibits relatively low stress values in the investigated operating regime. Accordingly, the paper's contribution lies in demonstrating that the proposed workflow can identify the most demanding feasible robot configurations, indicate whether substantially higher stresses are likely to remain hidden in unexplored regions of the configuration space, and support stress-aware trajectory screening during robot operation.

This manuscript is outlined as follows. Section 2 introduces the robot model and formulates worst-case stress identification as an optimisation problem in configuration space. Section 3 describes the configuration-stress data generation pipeline and scaling. Section 4 details the GP surrogate model, and Section 5 presents the BO procedure and experimental results. Section 6 discusses the interpretation, scope, and limitations of the proposed framework, and Section 7 concludes the paper.

2 Robot Model and Configuration-Space Problem Formulation

This section defines the mechanical setup and formulates the search for the stress-critical robot configuration as an optimisation problem. The manipulator considered here is shown in Figure 1 and detailed in [15], and its DH parameters are presented in Table 1.

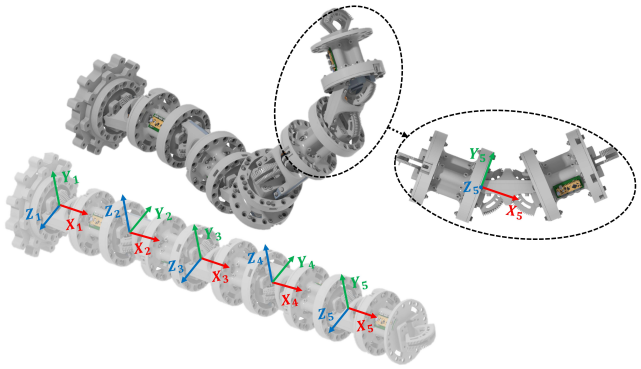


Fig. 1: Virtual prototype of the HR manipulator. The rendered picture shows the robot's local coordinate frames and a feasible configuration in space.

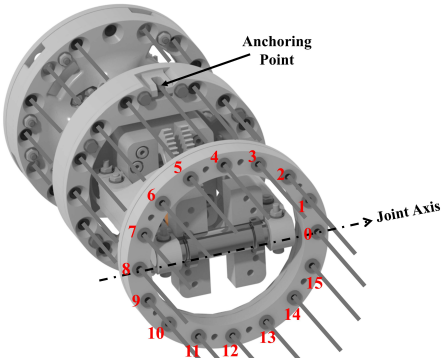


Fig. 2: Schematic of the cable routing ports distributed along the joint pitch circle.

For notational compactness, the independent joint variables listed in Table 1 are collected into the configuration vector

$$\mathbf{q} = [q_1, \ldots, q_{n_q}]^\top \in \mathcal{Q}, \tag{1}$$

where n_q is the number of independent configuration variables used in the optimisation.

2.1 Robot architecture and cable routing

The manipulator considered in this work is a five-joint cable-driven hyper-redundant robot. Each module has a nominal length of 0.1 m and a gear-based joint architecture, while the actuation is placed off-board and transmitted through antagonistic cable pairs. The routing geometry is therefore a central element of the mechanical model, as presented in [16], because each cable passes through several proximal modules before reaching its target joint, and the resulting cable directions and moment arms influence the internal load distribution. The cable-routing layout is encoded through a routing matrix

$$T_{\text{routing}} = \begin{bmatrix} 4 & 8 & 3 & 7 & 5 \\ 0 & 4 & 0 & 3 & 1 \\ 0 & 0 & 4 & 8 & 5 \\ 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 4 \end{bmatrix},$$

where each non-zero entry identifies the cable passage port associated with a given joint-cable interaction, as shown in figure 2. The pitch-circle diameter of the cable passages is

$$d_c = 0.115 \text{ m},$$

and the angular pitch between adjacent cable ports is

$$p_t = \frac{\pi}{8}.$$

From the routing matrix, the local cable-passage coordinates used in the statics model are obtained as

$$d_y = \frac{d_c}{2} \sin(T_{\text{routing}}\, p_t), \tag{2}$$

$$d_z = \frac{d_c}{2} \cos(T_{\text{routing}}\, p_t). \tag{3}$$

These quantities determine the radial directions and cable offsets that enter the local equilibrium equations of each module.

The geometric and inertial parameters used in the statics model are summarised in Table 2. They define the link lengths, cable pitch radius, local mass distribution, and cable preload values employed during the recursive equilibrium computation.

2.2 Configuration to reaction loads statics model

For each candidate configuration $\mathbf{q}$, the structural loading state is obtained through a recursive statics model of the cable-driven manipulator. The procedure first computes the forward kinematics and local reference frames of all joints, then solves static equilibrium sequentially from the distal end toward the base. In the present implementation, the external wrench at the end effector is set to zero,

$$\mathbf{F}_{\text{in}} = \mathbf{0}, \qquad \mathbf{M}_{\text{in}} = \mathbf{0},$$

so that the reported case study corresponds to self-weight, cable preload, and configuration-dependent internal loading without an additional external payload or contact load.

Let ${}^{w}R_j$ denote the rotation matrix from the j-th joint frame to the world frame. For each joint k, the local gravity vector is computed as

$$\mathbf{g}_k = ({}^{w}R_k)^\top \mathbf{g},$$

and the incoming force and moment from the distal chain are transformed into the local frame before solving the joint

Table 2: Main geometric, inertial, and preload parameters used in the configuration-to-reaction statics model.

Parameter	Value
Number of joints n_j	5
Module lengths l_1, l_2	0.100 m, 0.100 m
Cable pitch-circle diameter d_c	0.115 m
Cable pitch radius $r = d_c/2$	0.0575 m
Angular pitch p_t	$\pi/8$
d_{11} (all joints)	0.048 m
d_{21} (all joints)	0.0484 m
Link-side mass m_1 (all joints)	0.219 kg
Joint-side masses m_2	$[0.783,\ 0.771,\ 0.759,\ 0.779,\ 0.604]^\top$ kg
Cable preload matrix F_t	$\begin{bmatrix}100 & 100\\ 100 & 100\\ 100 & 100\\ 100 & 100\\ 100 & 100\end{bmatrix}$ N

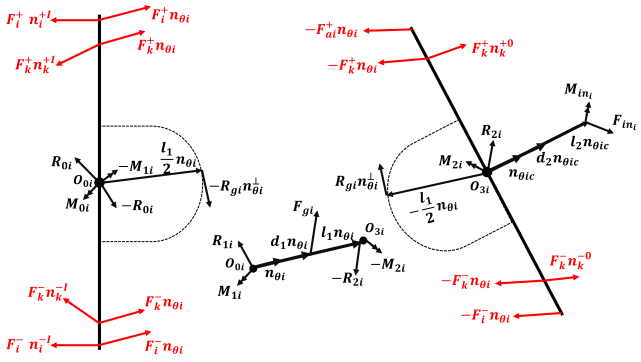


Fig. 3: Free body diagram of the joint module showing internal forces and cable interactions. Cable-joint interaction forces are shown in red.

equilibrium. At the joint level, the statics problem is written as a linear system of the form

$$A_k(\mathbf{q}_k, \mathcal{R}_k)\, \boldsymbol{\xi}_k = \mathbf{b}_k(\mathbf{F}_{\text{in},k}, \mathbf{M}_{\text{in},k}, \mathbf{g}_k, F_t),$$

where $\mathcal{R}_k$ denotes the local cable-routing geometry and

$$\boldsymbol{\xi}_k = \left[\mathbf{R}_{2,k}^\top\ \mathbf{M}_{2,k}^\top\ \mathbf{R}_{1,k}^\top\ \mathbf{M}_{1,k}^\top\ \mathbf{R}_{0,k}^\top\ \mathbf{M}_{0,k}^\top\ R_{g,k}\ \Delta F_{\text{cable},k}\right]^\top$$

collects the local reactions, local moments, a gravity-related internal term, and the cable-force redistribution variable. Reaction moments and forces, as well as the cable-joint interaction forces are showin in figure 3. The matrices A_k and $\mathbf{b}_k$ depend on the joint angle, the local cable passage directions, the cable pitch radius, the incoming wrench, and the preload state of the antagonistic cable pair. In practice, two equilibrium cases are solved depending on which cable of the antagonistic pair carries the higher load, so that the resulting cable forces satisfy the physically admissible tension condition at each joint.

Once the equilibrium of joint k is solved, the reactions transmitted to the previous module are propagated recursively according to

$$\mathbf{F}_{\text{in},k-1} = -\mathbf{R}_{0,k}, \qquad \mathbf{M}_{\text{in},k-1} = -\mathbf{M}_{0,k}. \tag{4}$$

Repeating this procedure from the distal module to the base yields the complete configuration-dependent loading state of the robot.

For the first-link component model considered in this paper, the loading transmitted from the remainder of the manipulator is represented through equivalent interface reactions extracted at the base of the recursive solution. The retained components are

$$\mathbf{x}(\mathbf{q}) = \left[M_x(\mathbf{q})\ M_y(\mathbf{q})\ R_x(\mathbf{q})\ R_y(\mathbf{q})\ R_z(\mathbf{q})\right]^\top, \tag{5}$$

where M_x and M_y are taken from the local joint moment vector $\mathbf{M}_{1,i}$, while R_x, R_y, and R_z are taken from the local reaction vector $\mathbf{R}_{1,i}$ for each joint. The torsional component about the local z-axis is negligible given the joint architecture, and is therefore not retained in the structural boundary conditions.

The adopted statics model is configuration-dependent and routing-dependent, but remains a quasi-static formulation. It does not explicitly model inertial effects (given the low moving speeds), distributed cable-hole friction along the full routing, or local contact compliance at the structural interfaces. Accordingly, the identified worst-case configuration should be interpreted as the worst case within the adopted static and component-level structural modelling assumptions.

2.3 Configuration parameterisation and worst-case stress objective

Each feasible manipulator configuration produces a unique internal loading state on the structure. For each candidate

configuration $\mathbf{q}$, the robot statics model is solved to compute the corresponding interface reactions and cable tensions. These quantities are then used to define the boundary conditions of the finite element model of the component of interest.

Let $\mathbf{x}(\mathbf{q})$ denote the vector of derived interface reactions obtained from the statics model:

$$\mathbf{x}(\mathbf{q}) = S(\mathbf{q}), \tag{6}$$

where $S(\cdot)$ denotes the configuration-to-reaction loads mapping. The component of interest is the first link of the robot. The scalar output of interest is the peak Von Mises stress generated in that component by the feasible configuration $\mathbf{q}$, denoted by

$$f(\mathbf{q}) = \sigma_{\mathrm{VM,max}}(\mathbf{q}). \tag{7}$$

The objective is therefore to identify the feasible robot configuration that maximises the peak stress:

$$\mathbf{q}^\star \in \arg\max_{\mathbf{q}\in Q} f(\mathbf{q}) = \arg\max_{\mathbf{q}\in Q} \sigma_{\mathrm{VM,max}}(\mathbf{q}). \tag{8}$$

Here $Q \subset \mathbb{R}^{n_q}$ denotes the feasible configuration set. In this formulation, the reaction vector $\mathbf{x}(\mathbf{q})$ is a configuration-dependent derived quantity used to define the structural boundary conditions; it is not the optimisation variable.

2.4 Black-box optimisation setting

Several factors make this an especially challenging optimisation task. First, the manipulator has a continuum of feasible poses, so the configuration space Q is continuous and potentially high-dimensional. Brute-force simulation over the feasible configuration space using Finite Element Analysis (FEA) is impractical. Furthermore, deriving a closed-form analytical model is insufficient for this application because simplified methods cannot reliably capture the complex geometry, interface conditions, local stress concentrations, and load-transfer mechanisms of the robot link. In the present study, high-fidelity FEA is required primarily to resolve geometric and interface-related stress concentrations rather than to model material nonlinearity, since all reported stress states remain within the linear elastic range of PA12. While FEA is necessary, simulating a very large configuration space is computationally prohibitive. This motivates the use of surrogate modelling, and in the present work, BO with GP is adopted because it is well suited to expensive black-box problems in low-data regimes and provides predictive uncertainty that can guide the search adaptively [17].

This setup is a classic expensive black-box optimisation problem: the mapping $\mathbf{q} \mapsto \sigma_{\mathrm{VM,max}}(\mathbf{q})$ is not known in closed form and can only be obtained by running a simulation. BO methods utilise a data-driven surrogate model of the objective and an acquisition strategy to select new sample points, thus balancing the exploration of unknown regions and the exploitation of high-performing regions [18]. Compared with non-adaptive exploration strategies such as uniform or random sampling, BO can use the available simulation budget more selectively by concentrating evaluations in regions that are either promising or uncertain [19]. In this application, a GP regression model is employed as the surrogate model for the stress function $f(\mathbf{q})$. The BO loop alternates between (i) fitting or updating the GP on available data, (ii) using an acquisition function (Expected Improvement) to propose the next configurations to evaluate, and (iii) validating those proposals with the FEA solver, which provides new stress data to refine the model. This process is shown in the schematic in figure 4. This closed loop repeats until a suitable stopping criterion is met, e.g. when the model's predictions are sufficiently accurate or the improvement becomes negligible. The ultimate objective is to converge to the maximum stress and identify the corresponding worst-case feasible robot configuration with high confidence in as few FEA calls as possible.

2.5 Configuration bounds and feasible domain

In the present implementation, the optimisation is carried out over a five-dimensional configuration domain defined by bounded joint variables,

$$Q = \left\{ \mathbf{q} \in \mathbb{R}^5 : q_j^{\min} \le q_j \le q_j^{\max}, \quad j = 1, \ldots, 5 \right\}. \tag{9}$$

For the reported case study, each configuration variable was bounded within

$$q_j \in [-45^\circ, 45^\circ], \qquad j = 1, \ldots, 5.$$

For each candidate configuration, the statics model is solved to compute the corresponding interface reactions used by the structural analysis. This guarantees that every evaluated point corresponds to a realisable robot configuration within the adopted configuration parameterisation.

3 Data Preparation and Scaling

Before fitting a machine-learning model, it is critical to preprocess the data into a suitable numerical form.

3.1 Initial configuration-stress dataset

The initial training dataset is defined as

$$\mathcal{D}_0 = \{(\mathbf{q}_i, y_i)\}_{i=1}^{n_0}, \tag{10}$$

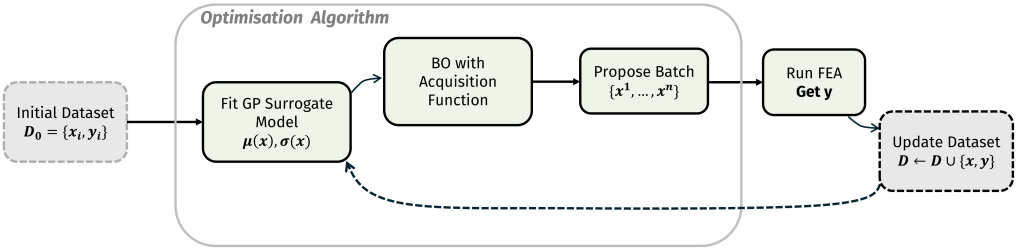


Fig. 4: Optimisation process overview

where $\mathbf{q}_i \in Q$ is a feasible robot configuration and

$$y_i = \sigma_{\mathrm{VM,max}}(\mathbf{q}_i) \tag{11}$$

is the corresponding peak Von Mises stress obtained by FEA. For each sampled configuration $\mathbf{q}_i$, the statics model is first solved to compute the derived interface reactions $\mathbf{x}(\mathbf{q}_i) = S(\mathbf{q}_i)$, which are then used to define the boundary conditions of the structural analysis. Thus, the surrogate model is trained directly on configuration-stress pairs, while the reaction vectors are treated as auxiliary outputs for engineering interpretation.

These raw data can have widely varying scales and units across the input dimensions, depending on the physical meaning and numerical range of the configuration variables. If left unscaled, features with large numerical ranges may dominate those with smaller ranges and adversely affect training. Likewise, the output stress values may include an arbitrary offset or a narrow range that slows convergence. For these reasons, a standard normalisation of inputs and outputs was performed prior to modelling.

3.2 Normalisation of inputs, outputs, and bounds

A per-feature z-score standardisation was applied to each input dimension, and a global standardisation to the output. Let $\mu_Q \in \mathbb{R}^{n_q}$ and $\sigma_Q \in \mathbb{R}^{n_q}$ be the mean and standard deviation of the configuration vectors in $\mathcal{D}_0$, and let μ_y and σ_y be the mean and standard deviation of the output stresses. The data are transformed as

$$\tilde{\mathbf{q}} = \mathrm{diag}(\sigma_Q)^{-1}(\mathbf{q} - \mu_Q), \tag{12}$$

$$\tilde{y} = \frac{y - \mu_y}{\sigma_y}. \tag{13}$$

This standardises every configuration variable and the output to zero mean and unit variance. All n_q inputs and the output are therefore placed on a common numerical scale.

The bounds of the search space are also transformed accordingly. If $\mathbf{q}^{\min}$ and $\mathbf{q}^{\max}$ are the physical lower and upper bounds of the configuration vector, then the corresponding normalised bounds are

$$\tilde{\mathbf{q}}^{\min} = \mathrm{diag}(\sigma_Q)^{-1}(\mathbf{q}^{\min} - \mu_Q), \tag{14}$$

$$\tilde{\mathbf{q}}^{\max} = \mathrm{diag}(\sigma_Q)^{-1}(\mathbf{q}^{\max} - \mu_Q). \tag{15}$$

This guarantees that the optimisation is performed on a properly scaled domain $[\tilde{\mathbf{q}}^{\min}, \tilde{\mathbf{q}}^{\max}] \subset \mathbb{R}^{n_q}$. After the GP model is trained and predictions are obtained in the normalised output scale, the results are always mapped back to physical units for interpretation.

4 Gaussian Process Surrogate Model

This section introduces the GP surrogate model used to approximate the stress response over the configuration space. The GP provides a probabilistic predictor (mean and uncertainty) in the normalised space, and we specify the prior (mean and kernel) and hyperparameter learning.

4.1 GP prior, mean function, and kernel

GPs are a flexible class of non-parametric models that define a distribution over functions, providing both predictions and uncertainty estimates [20]. It is assumed that the stress function

$$f : \mathbb{R}^{n_q} \to \mathbb{R}$$

is a sample from a GP prior, and that the observations are noisy samples of this function. Mathematically,

$$f \sim \mathcal{GP}(m, k), \qquad y = f(\mathbf{q}) + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, \sigma_n^2). \tag{16}$$

Here $m(\mathbf{q})$ is the mean function, i.e. the prior mean of f at input $\mathbf{q}$, and $k(\mathbf{q}, \mathbf{q}')$ is the covariance function (kernel), evaluated between two configuration vectors $\mathbf{q}$ and $\mathbf{q}'$ in the same input space. A constant mean $m(\mathbf{q}) = \beta$ is chosen, reflecting an unbiased prior with an unknown constant offset that will be learned from data. The kernel encodes prior assumptions about the function's smoothness and variation.

For this application, an Automatic Relevance Determination (ARD) Matérn 5/2 covariance function was used. The ARD Matérn-5/2 kernel is defined as [21]

$$k(\mathbf{q}, \mathbf{q}') = \sigma_f^2 \left(1 + \sqrt{5}\, r + \frac{5}{3} r^2\right) \exp(-\sqrt{5}\, r), \qquad r = \sqrt{\sum_{j=1}^{n_q} \frac{(q_j - q'_j)^2}{\lambda_j^2}}. \tag{17}$$

Here $k(\mathbf{q}, \mathbf{q}')$ is the covariance between the function values at two inputs $\mathbf{q}, \mathbf{q}' \in \mathbb{R}^{n_q}$, with $\mathbf{q}$ and $\mathbf{q}'$ denoting two (normalised) configuration vectors. The scalar σ_f^2 is the signal variance (overall output scale). The quantity r is the ARD-scaled Euclidean distance between $\mathbf{q}$ and $\mathbf{q}'$. The terms q_j and q'_j are the j-th components of $\mathbf{q}$ and $\mathbf{q}'$, respectively, λ_j is the length-scale associated with input dimension j, and $j \in \{1, \ldots, n_q\}$ indexes the configuration variables.

This kernel has the following hyperparameters: a length-scale λ_j for each input dimension and an overall signal variance σ_f^2. The sensitivity of the function to changes in input dimension j is controlled by λ_j: a small λ_j means $f(\mathbf{q})$ can vary rapidly, which implies that the corresponding configuration variable is highly relevant to predictions, whereas a large λ_j implies that $f(\mathbf{q})$ varies slowly with that variable. The Matérn-5/2 form is a popular choice in engineering design and BO because it produces sample functions that are twice differentiable, hence moderately smooth but not overly smooth [22]. In contrast, the ubiquitous squared-exponential (RBF) kernel assumes infinite differentiability, which may be too restrictive if the true stress response exhibits sharp changes. The Matérn family includes a parameter ν controlling smoothness; $\nu = 5/2$ gives a smoother process than Matérn-3/2 or Matérn-1/2, but less smooth than the RBF (which is equivalent to Matérn-$\nu \to \infty$). This makes the Matérn-5/2 kernel robust to modest fluctuations in the underlying function, which is appropriate here since the stress response may not be perfectly smooth over the configuration space. ARD further provides interpretability: after training, the relative magnitudes of λ_j indicate which configuration variables have the most significant effect on stress, with shorter length scales meaning greater relevance.

4.2 Hyperparameter learning via marginal likelihood

The Gaussian prior with the chosen $m(\mathbf{q})$ and $k(\mathbf{q}, \mathbf{q}')$ implies that, before seeing any data, any collection of function values

$$(f(\mathbf{q}^{(1)}), \ldots, f(\mathbf{q}^{(n)}))$$

has a joint multivariate normal distribution. Once observations are included, the GP posterior will also be Gaussian and can be obtained in closed form. The hyperparameters

$$\theta = \{\beta, \sigma_f^2, \lambda_1, \ldots, \lambda_{n_q}, \sigma_n^2\}$$

are not fixed a priori; they are determined from the data by maximising the Log Marginal Likelihood (LML) on the training set. The LML, also known as evidence, for observations $(\tilde{\mathbf{Q}}, \tilde{\mathbf{y}})$ in the normalised space is

$$L(\theta) = -\frac{1}{2}(\tilde{\mathbf{y}} - \beta\mathbf{1})^\top K_\theta^{-1}(\tilde{\mathbf{y}} - \beta\mathbf{1}) - \frac{1}{2}\log|K_\theta| - \frac{n}{2}\log(2\pi), \tag{18}$$

with

$$K_\theta = K_f + \sigma_n^2 I, \tag{19}$$

where K_f is the $n \times n$ kernel matrix of all training points, with entries

$$[K_f]_{ij} = k(\tilde{\mathbf{q}}_i, \tilde{\mathbf{q}}_j).$$

The function $L(\theta)$ is optimised with respect to θ to find the set of hyperparameters that best explains the data. This procedure is known as type-II maximum likelihood estimation or evidence maximisation for the GP [23]. The first term in Eq. (18) measures the model's ability to fit the training data. A better fit gives a higher value. The $\log|K_\theta|$ term acts as a penalty for making the model too complex. If the model is too flexible and tries to follow every small detail in the data, this term increases, discouraging overfitting. Thus, by maximising $L(\theta)$, the GP finds a good trade-off between accuracy and simplicity, fitting the data well without becoming too sensitive to noise. Once this optimisation is done, a fully trained GP model is obtained and can be used to make predictions.

4.3 Posterior prediction and uncertainty

Once the GP is trained, the predictive distribution at a new candidate configuration $\tilde{\mathbf{q}}_* \in \mathbb{R}^{n_q}$ is Gaussian with mean μ_* and variance σ_*^2. The posterior mean and variance are given by

$$\mu_* = \beta + \mathbf{k}_*^\top K_\theta^{-1}(\tilde{\mathbf{y}} - \beta\mathbf{1}), \tag{20}$$

$$\sigma_*^2 = k(\tilde{\mathbf{q}}_*, \tilde{\mathbf{q}}_*) - \mathbf{k}_*^\top K_\theta^{-1}\mathbf{k}_*. \tag{21}$$

Here $K_\theta \in \mathbb{R}^{n \times n}$ is the kernel matrix of the training inputs with noise on the diagonal,

$$[K_\theta]_{ij} = k(\tilde{\mathbf{q}}_i, \tilde{\mathbf{q}}_j) + \sigma_n^2 \delta_{ij},$$

$$\tilde{\mathbf{y}} = [\tilde{y}_1, \ldots, \tilde{y}_n]^\top,$$

and $\mathbf{k}_* \in \mathbb{R}^n$ is the kernel vector between each training point and the new point:

$$\mathbf{k}_* = [k(\tilde{\mathbf{q}}_1, \tilde{\mathbf{q}}_*), \ldots, k(\tilde{\mathbf{q}}_n, \tilde{\mathbf{q}}_*)]^\top. \tag{22}$$

Equation (20) shows that the predicted mean is a linear combination of observed residuals $(\tilde{y}_i - \beta)$, weighted by their covariance with $\mathbf{q}$; thus, predictions interpolate the data and extrapolate in between. Equation (21) shows that the predictive variance is the prior variance at $\mathbf{q}_*$ minus the amount explained by the data; it will be small near training points and larger in unexplored regions. As mentioned earlier, (μ, σ) is then converted back to physical units $(\hat{y}, \hat{\sigma})$. The GP not only provides an estimate of the peak stress at any feasible configuration, but also a quantitative uncertainty of that estimate. This uncertainty is essential for informing the BO, since it is useful to sample where the GP is unsure, as that region might hide a higher stress than currently predicted.

4.4 Computational considerations and model choice

It is worth noting the computational complexity of GP modelling. GP training and inference scale cubically $O(n^3)$ in the number of training points n, due to the inversion of the $n \times n$ kernel matrix, and quadratically in memory $O(n^2)$. In practice, this limits the number of exact GPs to a few thousand data points on a single machine before computations become impractical. In the present structural stress application, the number of affordable FEA samples is relatively small, initial data plus iterative additions, usually with n in the low hundreds at most, so an exact GP is feasible. This surrogate forms the core of the Bayesian optimisation loop, because it allows the acquisition function to be computed without actually running the FEA at every candidate point. If the true stress function had strong non-stationarities or varying smoothness in different regions, more advanced models such as Deep GPs, i.e. multiple GP layers that can learn input warpings, could be considered [24].

5 Bayesian Optimisation

This section outlines the BO loop used to efficiently search the feasible configuration space for worst-case (maximum) stress using the GP surrogate. We define the acquisition function that trades off exploration and exploitation, and describe how new simulation points are selected and incorporated.

5.1 Acquisition function: Expected Improvement

With a trained GP model for $f(\mathbf{q})$, BO proceeds by repeatedly asking: "Where should we sample the FEA simulation next to discover higher stress?" BO addresses this by using an acquisition function $\alpha(\mathbf{q})$ that quantifies the utility of sampling any given configuration $\mathbf{q}$, based on the model's current predictions. The acquisition function is designed to balance exploration, by sampling regions where model uncertainty is high, and exploitation, by sampling regions where the model's mean prediction is high. This balance is important: if it only exploits, it risks getting stuck in a local optimum and missing the true worst-case scenario; if it only explores, evaluations are wasted in areas that may not induce higher stress than what was already observed.

For the maximisation problem, the Expected Improvement (EI) acquisition function was employed, a widely used choice in BO [25]. At any iteration, let f_{best} be the current best observed value of the stress in the normalised scale, i.e. the maximum $\tilde{y}$ among all evaluated points so far. For a candidate configuration $\mathbf{q}$, the one-step improvement is defined as $(f(\mathbf{q}) - f_{\text{best}})^+$, where the positive-part operator $(\cdot)^+$ is given by $z^+ = \max(z, 0)$ for any real number z. Since $f(\mathbf{q})$ is a random variable under the GP, the EI can be written in closed form as

$$\text{EI}(\mathbf{q}) = \mathbb{E}\left[(f(\mathbf{q}) - f_{\text{best}})^+\right] = (\mu_* - f_{\text{best}})\Phi(Z) + \sigma_*\phi(Z) \tag{23}$$

where

$$Z = \frac{\mu_* - f_{\text{best}}}{\sigma_*}, \tag{24}$$

and $\phi(Z)$ and $\Phi(Z)$ are the standard normal probability density and cumulative distribution functions, respectively. Here $\mu_* = \mathbb{E}[f(\mathbf{q})]$ and $\sigma_* = \text{Std}[f(\mathbf{q})]$ are the GP predictive mean and standard deviation at $\mathbf{q}$. EI increases when the predicted mean is high or when uncertainty is high, thus balancing sampling in promising regions and reducing uncertainty in unexplored regions. The term $(\mu_* - f_{\text{best}})\Phi(Z)$ captures the contribution of the predicted mean, while the term $\sigma_*\phi(Z)$ captures the contribution of uncertainty. As a result, EI favours configurations with either high predicted stress (exploitation) or high uncertainty (exploration), providing an automatic trade-off between the two without requiring an explicit exploration parameter.

5.2 Batch Bayesian optimisation with batch Expected Improvement

To identify the worst-case stress configuration, a batch BO strategy was employed using a GP model and an EI acquisition function. Rather than querying one configuration at a time, a batch of B points

$$Q_{\text{batch}} = \{\mathbf{q}^{(1)}, \mathbf{q}^{(2)}, \ldots, \mathbf{q}^{(B)}\}$$

was selected to evaluate in parallel. In the experiments reported here, the batch size was set to $B = 10$. The batch Expected Improvement acquisition function is defined as the

EI obtained by evaluating all B points collectively. Mathematically, for a batch Q_{batch},

$$\alpha_{\text{BEI}}(Q_{\text{batch}}) = \mathbb{E}\left[\max\left(0, \max_{i=1,\ldots,B}\left(f(\mathbf{q}^{(i)}) - f_{\text{best}}\right)\right)\right]. \tag{25}$$

This is the expected increase in the best-observed stress after simultaneously sampling the set Q_{batch}. The batch-EI criterion evaluates the joint distribution of $(f(\mathbf{q}^{(1)}), \ldots, f(\mathbf{q}^{(B)}))$ under the GP, so that the B points are chosen together to maximise their collective EI. In effect, the acquisition function accounts for correlations between points in the batch and rewards diversity while targeting high-stress regions of the input space. All B proposed robot configurations are evaluated in parallel, generating B new observations per iteration.

Computing batch-EI for $B > 1$ is challenging because there is no closed-form expression for the expectation of the maximum of B correlated outcomes. In the current implementation, a Monte Carlo simulation-based estimator was employed to approximate $\alpha_{\text{BEI}}(Q_{\text{batch}})$. Essentially, the GP posterior is repeatedly sampled to estimate the EI for a given batch. An inner optimisation loop, using a multi-start gradient-based method, is then employed to find the optimal batch $Q^{\star}_{\text{batch}}$ that maximises the batch-EI acquisition over the domain. This produces B promising configurations to evaluate next. These are chosen so that each point targets a high-predicted-stress region, an uncertain region, or both.

After proposing a batch, all B configurations are evaluated through the FEA solver, which automatically receives the output from the BO algorithm and simulates the corresponding structural cases using the method described in [26]. The resulting stress values $y^{(1)}, \ldots, y^{(B)}$ are the true outputs for those inputs. These new data points are then added to the dataset, and the GP model is updated to fit the new observations. In this closed-loop manner, the BO alternates between surrogate modelling, acquisition optimisation, and high-fidelity evaluation on the FEA solver. With each iteration, the GP model becomes more accurate in the high-stress regions and the value of f_{best} increases. This process continues until a convergence criterion is met. In the present implementation, termination is based on two complementary indicators: (i) the maximum Expected Improvement (EI) associated with the proposed batch, and (ii) the relative improvement in the best observed stress value. More specifically, the algorithm is stopped only when both the normalised EI and the relative increase in the best observed stress fall below prescribed thresholds for three consecutive iterations. This dual criterion is intended to indicate both acquisition saturation and negligible progress in the objective, thereby reducing the likelihood of stopping while potentially critical unexplored regions still remain. At termination, the algorithm returns the highest observed stress and the corresponding feasible robot configuration as the solution.

5.3 Model performance evaluation

The purpose of the following evaluation is not to establish a universal ranking among optimisation algorithms or surrogate models. Instead, the objective is to verify that the proposed BO-GP workflow is practically useful for this engineering problem: namely, that it can learn the configuration-stress relationship with sufficient accuracy, guide the search toward stress-critical feasible configurations, and support worst-case structural assessment under a limited FEA budget.

The model was trained and evaluated for the first link of the HR manipulator. Since this prototype was designed to withstand future higher stresses from the addition of modules, the links are significantly overdesigned. For this reason, the stress values on this component are relatively low. Nevertheless, the first link serves as a suitable validation platform for the proposed optimisation method.

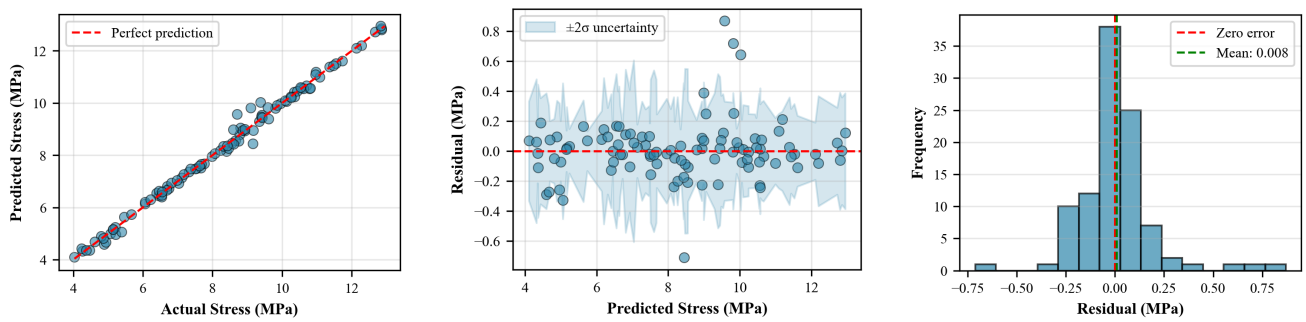
The reported implementation uses an initial design of $N_0 = 200$ feasible robot configurations, batch proposals of size $B = 10$, and a convergence check based on both acquisition saturation and diminishing improvement in the best observed stress. The optimisation terminated after 6 BO iterations, corresponding to 250 total evaluated configurations. The final model performance was assessed on a separate test set of 100 previously unseen feasible configurations.

Over the course of the BO iterations, the framework progressively identified higher-stress regions of the configuration space and converged to a stress-critical feasible configuration

$$\mathbf{q}^{\star} = [0.7854,\ 0.5781,\ -0.7854,\ 0.0170,\ -0.0304]^{\top}\ \text{rad},$$

with the corresponding peak stress $\sigma(\mathbf{q}^{\star}) = 14.23$ MPa. Relative to the best configuration contained in the initial design, this corresponds to an absolute increase of 0.20 MPa, i.e. a relative improvement of 1.41%. The relatively small improvement over the best configuration already present in the initial design suggests that the initial sample set had already captured part of the high-stress region, while the BO iterations further refined the search and confirmed the most demanding feasible configuration within the adopted domain. The corresponding derived reaction vector $\mathbf{x}(\mathbf{q}^{\star})$ can be reported as an auxiliary output of the statics model for engineering interpretation.

Figure 5a shows the relation between the predicted and actual peak stress compared with the ideal one-to-one trend, while Figure 5b reports the residuals and reveals whether a systematic bias is present across the stress range. The residuals remain centred close to zero, with a mean residual of 0.010 MPa and a standard deviation of 0.204 MPa, indicating that the surrogate reproduces the FEA response without



(a) Predicted vs. Actual VM Stress (b) Residuals vs. FEA Actual Stress (c) Distribution of Residuals

Fig. 5: Predictive performance of the configuration-stress surrogate model on 100 previously unseen feasible robot configurations: (a) predicted versus actual peak stress, (b) residuals versus actual stress, and (c) residual histogram.

a strong systematic offset. All reported stresses lie within the linear elastic range of PA12, so the observed residual behaviour is not attributed to material nonlinearity. Figure 5c presents the histogram of the residuals over the test set; the concentration around zero and the scarcity of large-magnitude residuals suggest that substantial prediction errors are not common in the analysed regime. This behaviour is consistent with the BO/EI sampling strategy, which progressively concentrates evaluations in regions that are either uncertain or likely to contain high stress, thereby improving local model accuracy near critical configurations.

The Relative Error (ϵ), the Root Mean Squared Error (RMSE), the Mean Absolute Error (MAE), and the Determination Coefficient (R^2) are calculated as

$$\varepsilon_{\text{rel}} = \frac{|y_{\text{pred},i} - y_{\text{test},i}|}{y_{\text{test},i}} \times 100\%. \tag{26}$$

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(y_{\text{test},i} - y_{\text{pred},i}\right)^2}. \tag{27}$$

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}\left|y_{\text{test},i} - y_{\text{pred},i}\right|. \tag{28}$$

$$R^2 = 1 - \frac{\sum_{i=1}^{n}\left(y_{\text{test},i} - y_{\text{pred},i}\right)^2}{\sum_{i=1}^{n}\left(y_{\text{test},i} - \bar{y}_{\text{test}}\right)^2}. \tag{29}$$

Here n is the total number of samples, $y_{\text{test},i}$ is the actual observed value for the i-th point in the test set, $y_{\text{pred},i}$ is the predicted value for the i-th data point in the test set, and $\bar{y}_{\text{test}}$ is the mean of the actual values in the test set. The obtained metrics confirm that the surrogate is sufficiently accurate to guide the FEA-based search toward stress-critical feasible configurations within the simulation budget adopted in this study, as summarised in Table 3.

Table 3: Summary of dataset composition, preprocessing parameters, and GP performance metrics.

Metric / Parameter	**Value**
Dataset and Preprocessing	
Input Features	5
Initial Training Samples N_0	200
Batch Size B	10
BO Iterations	6
Total Evaluated Samples	250
Test Samples	100
Standardization Method	Standard Scaler
Output Mean (before scaling)	8.02 MPa
Output Std Deviation (before scaling)	0.07 MPa
Model Performance (Test Set)	
MAE	0.126 MPa
RMSE	0.205 MPa
Mean ϵ	1.61%
Maximum ϵ	12.71%
R^2	0.992
Mean Residual	0.010 MPa
Residual Std Dev	0.204 MPa
Mean Predicted Std Dev	0.161 MPa

Table 4 presents the GP model configuration parameters and hyperparameters, all reported in the normalised input space defined by Eq. (12).

The learned hyperparameters provide additional information about the surrogate model's behaviour across the sampled configuration space. In particular, the ARD length scales offer a useful indication of the relative sensitivity of the peak stress to changes in the different configuration variables, thereby supporting the interpretation of which joint coordinates have the strongest influence within the analysed operating regime.

An additional consistency check was performed by optimising the final GP surrogate after convergence and comparing its predicted maximum with the best stress value actually

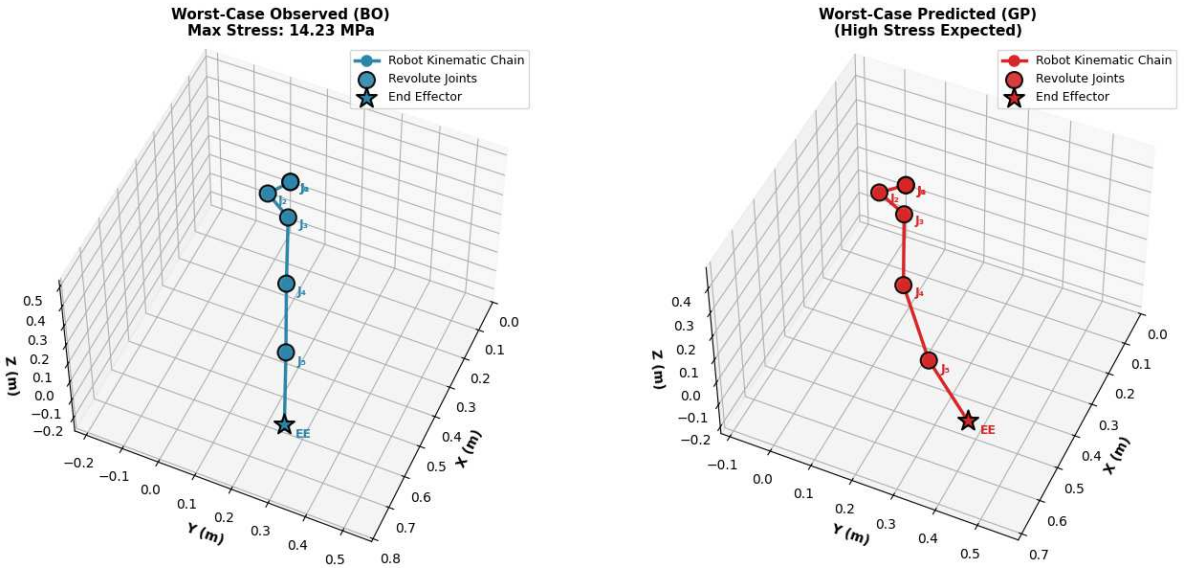


Fig. 6: Comparison between the worst-case configuration directly observed during the BO-FEA loop (left) and the high-stress configuration predicted by the final GP surrogate (right).

Table 4: GP model configuration and hyperparameters.

Component	**Value**
Noise Variance (σ_n^2)	6.93×10^{-1}
Output Scale	2.19×10^{-1}
Mean Constant	0.39
Lengthscales (per configuration variable in normalised space)	
Configuration variables q_1 to q_5	0.76, 1.83, 2.33, 5.16, 5.01
Lengthscale Mean	3.02
Lengthscale Std Dev	1.76
MLL Optimizer	L-BFGS-B (BoTorch)

observed during the BO-FEA loop. The GP-predicted high-stress configuration was

$$\mathbf{q}_{\mathrm{GP}} = [0.7854,\ 0.6184,\ -0.7854,\ 0.3079,\ 0.3617]^{\top}\ \mathrm{rad},$$

for which the final surrogate predicted a peak stress of 14.06 MPa. Relative to the best-observed BO configuration, this corresponds to a surrogate gap of −0.17 MPa, i.e. −1.22%. When the same configuration was subsequently re-evaluated in ANSYS, it produced approximately 14.2 MPa.

Figure 6 provides a geometric comparison between the worst-case configuration directly observed during the BO-FEA loop and the high-stress configuration predicted by the final GP surrogate. Although the two robot poses are visibly different, especially in the distal portion of the chain, the corresponding peak stress values are very close. In particular, the GP-predicted configuration was subsequently validated in ANSYS and produced a peak stress of approximately 14.2 MPa, which is nearly identical to the 14.23 MPa maximum observed during BO.

6 Discussion

6.1 Interpretation of the low-stress case study

The present case study should be interpreted in light of the analysed component and the adopted operating conditions. The first link considered here is intentionally overdesigned and therefore does not approach a material failure threshold within the investigated configuration domain. Accordingly, the absolute maximum stress identified in this study, 14.23 MPa, is not by itself the main result of interest. Rather, the significance of the case study is methodological: it shows that the proposed workflow can identify the most demanding feasible configuration for the analysed component and determine whether substantially higher stresses are likely to remain hidden in unexplored regions of the adopted configuration space.

From this perspective, the value of the framework is not limited to locating a single worst-case sample. Instead, it provides a practical basis for design verification and for stress-dependent operational screening. In particular, even when the analysed component is overdesigned, the same configuration-space BO-GP formulation can be used to identify robot poses associated with elevated structural demand and to exclude such poses from motion planning or experimental testing. This use case is particularly relevant for practical robot operation, since the same surrogate-assisted configuration-space formulation can be employed not only to identify worst-case poses at design time, but also to define safe operating regions by screening out configurations associated with elevated structural stress.

6.2 Convergence, consistency, and critical stress regions

The optimisation results indicate that the initial set of 200 feasible configurations already contained samples in the vicinity of the high-stress region, since the best stress improved by only 1.41% during the BO stage. This should not be interpreted as a weakness of the method. On the contrary, it shows that the BO iterations refined and confirmed the most demanding region of the adopted search domain rather than drifting toward unrelated regions of configuration space.

A particularly important result is that the final GP surrogate identified a geometrically distinct configuration,

$$\mathbf{q}_{\mathrm{GP}} = [0.7854,\ 0.6184,\ -0.7854,\ 0.3079,\ 0.3617]^{\top}\ \mathrm{rad},$$

for which the surrogate predicted a peak stress of 14.06 MPa, while a subsequent ANSYS validation returned approximately 14.2 MPa. Although this configuration differs visibly from the BO-observed worst-case pose, the resulting stress is nearly identical to the 14.23 MPa maximum observed during the BO-FEA loop. This suggests that the stress-critical condition is not represented by a single isolated optimum, but rather by a neighbourhood of feasible configurations that generate similarly high stresses.

This interpretation is reinforced by the small surrogate gap of -1.22% between the best observed stress and the best surrogate-predicted maximum at convergence. Taken together, the BO optimum, the GP-predicted near-critical alternative, and the ANSYS re-evaluation suggest that the workflow did not terminate while leaving a markedly more demanding configuration unresolved within the adopted search domain. From an engineering standpoint, this is valuable because real robot motion is better characterised by identifying classes of dangerous poses or high-stress regions than by a single isolated optimum alone.

6.3 Scope and limitations

The present case study is restricted to the first link and should therefore be interpreted as a component-level demonstration of the framework rather than a complete structural assessment of the robot. The structural analysis is driven by equivalent interface reactions obtained from the recursive statics model, so the identified optimum corresponds to the worst case within the adopted static and component-level modelling assumptions. In the reported implementation, the end-effector wrench is set to zero in the recursive statics model, so the analysed load cases arise from self-weight, cable preload, and configuration-dependent internal load transfer rather than from additional external payload or contact loads.

In other operating scenarios, especially when external loads act not only at the end effector but also at intermediate sections, the location of the maximum stress may shift to other links depending on their geometry, connection details, and load-transfer characteristics. For this reason, the current results should not be generalised to the whole manipulator without further analysis. Extending the same configuration-space workflow to multiple links, or to a robot-level objective defined as the maximum stress over all links, is therefore an important direction for future work.

7 Conclusions

This paper presented a computationally efficient workflow for identifying stress-critical feasible robot configurations for a manipulator component using a limited number of finite element evaluations. The key output of the method is the best-observed stress value and its associated feasible configuration $\mathbf{q}^{\star}$, which provides a practical reference for structural verification within the admissible operating range of the robot. In the reported case study, the BO-GP workflow was initialised with 200 feasible configurations and converged after 6 optimisation iterations, corresponding to 250 total evaluated configurations. The identified stress-critical configuration produced a peak stress of 14.23 MPa, while the learned surrogate achieved a mean absolute error of 0.13 MPa and an R^2 value of 0.992 on a separate test set of 100 unseen configurations. These results show that the proposed framework can serve as a practical tool for configuration-space structural assessment in a cable-driven hyper-redundant robot under constrained FEA limitations.

The purpose of the proposed framework is not to establish the superiority of this optimisation algorithm or surrogate model over others. Rather, the contribution lies in demonstrating that a BO-GP workflow can be used as a physically consistent engineering tool for worst-case configuration identification, structural verification, and safe operating-range evaluation.

This optimisation framework focuses on peak Von Mises stress as the scalar objective; therefore, other failure modes such as fatigue, wear, or local contact damage are outside the scope of this work. Future developments may extend the same configuration-space framework to additional robot links, alternative structural performance metrics, or constrained multi-objective formulations.

Data Availability

Further data supporting the findings of this study and allowing reproduction of the results are made available at [27].

Declarations

Ethics approval and consent to participate:Not applicable.

Consent for publication: Not applicable.

Funding:This research received no external funding

Authors' contributions: Ardit Poka - Writing, Methodology, Analysis, Data curation, Coding, Simulation, Conceptualisation;
Daniele Ludovico - Review and editing, Validation, Coding;
Lorenzo De Mari Casareto Dal Verme - Review and editing;
Federico Manara - Review & editing;
Carlo Canali - Supervision, Resources, Project administration, Investigation, Funding acquisition;
Giovanni Berselli - Supervision, Conceptualisation, Investigation;
Darwin G. Caldwell - Supervision, funding acquisition.

Acknowledgements: Not applicable

Conflict of Interest: The authors have no relevant financial or non-financial interests to disclose.

References

1. Mengesha, G., 2025. "Design optimization in structural engineering: A systematic review of computational techniques and real-world applications". *Available at SSRN 5254589*.
2. Prayogo, D., Hauwing, W., Wong, F., Tjandra, D., et al., 2020. "Structural design optimization using particle swarm optimization and its variants". In IOP Conference Series: Earth and Environmental Science, Vol. 506, IOP Publishing, p. 012048.
3. Karafi, Y., Moussaoui, Z., and Abou El Majd, B., 2024. "Simultaneous surrogate modeling and dimension reduction using unsupervised learning. application to parametric wing shape optimization". *Math. Model. Comput,* **11**(1), pp. 154–165.
4. Tromme, E., Tortorelli, D., Brüls, O., and Duysinx, P., 2015. "Structural optimization of multibody system components described using level set techniques". *Structural and Multidisciplinary Optimization,* **52**(5), pp. 959–971.
5. Pistone, A., Ludovico, D., Dal Verme, L. D. M. C., Leggieri, S., Canali, C., and Caldwell, D. G., 2024. "Modelling and control of manipulators for inspection and maintenance in challenging environments: A literature review". *Annual Reviews in Control,* **57**, p. 100949.
6. Manara, F., Poka, A., Ludovico, D., Pistone, A., Dal Verme, L. D. M. C., Canali, C., Berselli, G., and Caldwell, D. G., 2025. "Tendon-driven hyper-redundant manipulators: A review on design solutions". *IEEE/ASME Transactions on Mechatronics*.
7. Poka, A., Manara, F., Ludovico, D., Pistone, A., Dal Verme, L. D. M. C., Canali, C., Berselli, G., and Caldwell, D. G., 2025. "Alternative solutions to tendon-driven hyper-redundant manipulators: A review on design and actuation strategies". *IEEE/ASME Transactions on Mechatronics*.
8. Canali, C., Pistone, A., Ludovico, D., Guardiani, P., Gagliardi, R., De Mari Casareto Dal Verme, L., Sofia, G., and Caldwell, D. G., 2022. "Design of a novel long-reach cable-driven hyper-redundant snake-like manipulator for inspection and maintenance". *Applied Sciences,* **12**(7), p. 3348.
9. Wang, X., Jin, Y., Schmitt, S., and Olhofer, M., 2023. "Recent advances in bayesian optimization". *ACM computing surveys,* **55**(13s), pp. 1–36.
10. Paulson, J. A., and Tsay, C., 2025. "Bayesian optimization as a flexible and efficient design framework for sustainable process systems". *Current Opinion in Green and Sustainable Chemistry,* **51**, p. 100983.
11. Tran, Q. D., Shin, D., and Jang, G.-W., 2024. "Bayesian optimization-based topology optimization using moving morphable bars for flexible structure design problems". *Engineering Structures,* **300**, p. 117103.
12. Kim, Y., Pan, Z., and Hauser, K., 2021. "Mo-bbo: Multi-objective bilevel bayesian optimization for robot and behavior co-design". In 2021 IEEE International Conference on Robotics and Automation (ICRA), IEEE, pp. 9877–9883.
13. Jenkins, W. F., Gerstoft, P., and Park, Y., 2023. "Bayesian optimization with gaussian process surrogate model for source localization". *The Journal of the Acoustical Society of America,* **154**(3), pp. 1459–1470.
14. Qin, N., Zhou, X., Wang, J., and Shen, C., 2021. "Bayesian optimization: Model comparison with different benchmark functions". In 2021 International Conference on Signal Processing and Machine Learning (CONF-SPML), IEEE, pp. 329–333.
15. Poka, A., 2026. "Design, optimisation, and actuation of tendon-driven hyper-redundant manipulators".
16. Guardiani, P., Ludovico, D., Pistone, A., Abidi, H., Zaplana, I., Lee, J., Caldwell, D. G., and Canali, C., 2022. "Design and analysis of a fully actuated cable-driven joint for hyper-redundant robots with optimal cable routing". *Journal of Mechanisms and Robotics,* **14**(2), p. 021006.
17. Rajaram, D., Puranik, T. G., Ashwin Renganathan, S., Sung, W., Fischer, O. P., Mavris, D. N., and Ramamurthy, A., 2021. "Empirical assessment of deep gaussian process surrogate models for engineering problems". *Journal of Aircraft,* **58**(1), pp. 182–196.
18. Candelieri, A., 2021. "A gentle introduction to bayesian optimization". In 2021 Winter Simulation Conference (WSC), IEEE, pp. 1–16.
19. Diessner, M., O'Connor, J., Wynn, A., Laizet, S., Guan, Y., Wilson, K., and Whalley, R. D., 2022. "Investigating bayesian optimization for expensive-to-evaluate black box functions: Application in fluid dynamics". *Frontiers in Applied Mathematics and Statistics,* **8**, p. 1076296.
20. Liu, H., Ong, Y.-S., Shen, X., and Cai, J., 2020. "When gaussian process meets big data: A review of scalable gps". *IEEE transactions on neural networks and learning systems,* **31**(11), pp. 4405–4423.
21. Grigorievskiy, A., and Karhunen, J., 2016. "Gaussian process kernels for popular state-space time series models". In 2016 International Joint Conference on Neural Networks (IJCNN), IEEE, pp. 3354–3363.
22. Le Riche, R., and Picheny, V., 2021. "Revisiting bayesian optimization in the light of the coco benchmark". *Structural and Multidisciplinary Optimization,* **64**(5), pp. 3063–3087.
23. Butler, A., Haynes, R. D., Humphries, T. D., and Ranjan, P., 2014. "Efficient optimization of the likelihood function in gaussian process modelling". *Computational Statistics & Data Analysis,* **73**, pp. 40–52.
24. Damianou, A., and Lawrence, N. D., 2013. "Deep gaussian processes". In Artificial intelligence and statistics, PMLR, pp. 207–215.
25. Gan, W., Ji, Z., and Liang, Y., 2021. "Acquisition functions in bayesian optimization". In 2021 2nd international conference on big data & artificial intelligence & software engineering (ICBASE), IEEE, pp. 129–135.
26. Poka, A., Ludovico, D., Manara, F., De Mari Casareto Dal Verme, L., Canali, C., Berselli, G., and Caldwell, D. G., 2025. "Automation of the structural analysis on a robot link via scripting in ansys

workbench: A tutorial". In International Conference of the Italian Association of Design Methods and Tools for Industrial Engineering, Springer, pp. 269–280.
27. Ardit Poka, 2026. https://github.com/iit-rain/worst-case-stress-identification-on-a-hyper-redundant-robot-using-bayesian-optimisation.