

AI-Assisted Design Optimization of Modular Robotic Arm Structures Using Genetic- Algorithm-Trained Neural Surrogates

Timileyin Opeyemi Akande

akande.timileyin@tech-u.edu.ng

Abiola Ajimobi Technical University Ibadan

Research Article

Keywords: genetic algorithm, neural network surrogate, modular robotic arm, structural design optimization, NSGA-II, finite element analysis, lightweight design, multi-objective optimization

Posted Date: October 27th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11085266/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

AI-Assisted Design Optimization of Modular Robotic Arm Structures Using Genetic-Algorithm-Trained Neural Surrogates

Timileyin Opeyemi AKANDE^{1*}  0000-0003-4307-1880

¹Department of Mechanical and Mechatronics Engineering, Abiola Ajimobi Technical University (First Technical University), Ibadan, Oyo State, 200255, Nigeria

*Corresponding author's email: akande.timileyin@tech-u.edu.ng

Abstract

Modular robotic arms are increasingly adopted in manufacturing and service robotics because interchangeable link and joint modules shorten redesign cycles and improve part commonality across product families. Sizing each module, however, still normally requires repeated finite-element analysis (FEA), which becomes computationally prohibitive once it is embedded inside a population-based search such as a genetic algorithm (GA). This paper presents a genetic-algorithm-trained neural-surrogate framework that replaces the great majority of these expensive structural evaluations with a fast, learned approximation. A two-module shoulder-to-elbow and elbow-to-wrist link chain is parameterized by six design variables covering outer diameter, wall thickness, and material selection per module. A Latin-Hypercube-sampled dataset of 1,500 configurations is labelled using a one-dimensional Euler-Bernoulli finite-element beam model that returns structural mass, tip deflection, root and step-section von Mises stress, and first natural frequency. A feed-forward neural network trained on this dataset attains coefficients of determination between 0.984 and 0.997 on an independent test set. The trained surrogate is then embedded as the fitness evaluator inside a constrained non-dominated sorting genetic algorithm (NSGA-II) that searches for Pareto-optimal trade-offs between mass and tip deflection subject to stress, frequency, and buckling constraints. The surrogate-assisted optimizer converges to the same optimum as a direct finite-element-in-the-loop genetic algorithm while consuming roughly one-eighth of the high-fidelity structural evaluations, and 98.8% of its Pareto-optimal candidates remain feasible under independent re-validation. The best surrogate-optimized configuration reduces structural mass by 19.9% relative to a representative baseline while satisfying every specified constraint, demonstrating a reproducible and computationally economical route to modular robotic-arm structural design.

Keywords: *genetic algorithm; neural network surrogate; modular robotic arm; structural design optimization; NSGA-II; finite element analysis; lightweight design; multi-objective optimization*

1. Introduction

Modular design has become a central strategy in robotic-arm engineering because interchangeable joint and link modules allow a single hardware catalogue to be reconfigured into many task-specific kinematic chains without redesigning the manipulator from first principles [1]. This approach shortens development cycles, lowers manufacturing cost through part commonality, and simplifies maintenance, and it has motivated a substantial body of recent work on reconfigurable and modular manipulator architectures [1]. Digital-twin and Industry 4.0 initiatives have further increased the pressure to make the structural design of each module traceable, data-driven, and rapidly re-optimizable whenever a manipulator is reconfigured for a new task [2]. Yet however flexible the assembly-level architecture becomes, every module in the library must still be individually sized: its cross-sectional geometry, wall thickness, and material must be chosen so that the assembled arm

satisfies mass, stiffness, strength, and vibration requirements across the range of configurations and payloads in which that module will be used.

Structural sizing of robotic-arm links has traditionally relied on iterative finite-element analysis (FEA), in which a candidate geometry is meshed, solved, and post-processed for stress, deflection, and modal response before being accepted or revised [3]–[10]. When this evaluation loop is embedded inside a population-based search such as a genetic algorithm, the number of expensive FEA calls scales with the product of population size and generation count, and published robot-arm optimization studies routinely couple a GA directly to FEA or an equivalent physics solver [11], [6]. For a single fixed-configuration arm this cost is manageable within a single design cycle; for a modular arm whose links must be re-optimized across several configurations, payloads, and candidate materials, the cumulative simulation burden discourages designers from exploring the full trade-off space between competing objectives such as mass and stiffness.

Machine-learning surrogate models offer a practical way to relieve this bottleneck. A neural network trained on a modest number of high-fidelity simulations can approximate the simulator's input-output mapping to high accuracy and then evaluate new candidate designs in a fraction of a millisecond [12]–[18]. Embedding such a surrogate inside an evolutionary optimizer, in place of the simulator itself, allows the search to explore orders of magnitude more candidate designs for the same computational budget a strategy that has been demonstrated for topology-optimized components [12], [13], impact and dynamic-response prediction [14], [15], one-dimensional structural systems [17], architected metamaterials [18], and buckling-critical lightweight structures optimized with NSGA-II [16]. Surrogate-assisted evolutionary algorithms (SAEAs) have consequently become an active research area in their own right [19], [20]. However, the specific combination that this paper addresses a modular robotic-arm link chain in which both continuous geometric variables and discrete material choice are optimized simultaneously by a genetic algorithm trained on, and validated against, an independent structural surrogate has not been reported as a single reproducible pipeline in the literature reviewed in Section 2.

This paper makes the following contributions:

- 1) A parametric two-module robotic-arm link-chain model together with a reduced-order one-dimensional finite-element (FE) beam evaluator that plays the role of the “ground-truth” high-fidelity simulator, generating labelled structural-response data (mass, tip deflection, root and step stress, first natural frequency) fast enough to build a training set of thousands of samples.
- 2) A feed-forward neural-network surrogate trained on this dataset that predicts the three simulation-expensive responses (deflection, stress, natural frequency) with coefficients of determination between 0.984 and 0.997 on an independent, differently-seeded test set.
- 3) A from-scratch, constrained implementation of the non-dominated sorting genetic algorithm (NSGA-II) that uses the trained surrogate as its in-the-loop fitness evaluator, together with a matched direct FE-in-the-loop NSGA-II run used as a control experiment.
- 4) A quantitative demonstration that the surrogate-assisted optimizer converges to the same Pareto front as the direct FE-in-the-loop optimizer while consuming markedly fewer high-fidelity

evaluations, together with an explicit re-validation step showing that 98.8% of the surrogate-optimized Pareto set remains feasible when re-evaluated with the ground-truth model.

- 5) Design-level findings for the case-study arm, including a consistent preference for carbon-fibre-reinforced-polymer (CFRP) modules across the entire Pareto front and a quantified sensitivity ranking of the six design variables.

The remainder of the paper is organized as follows. Section 2 reviews related work on robotic-arm structural optimization, genetic algorithms, and neural-network surrogates. Section 3 formulates the modular-arm sizing problem mathematically and then details the proposed genetic-algorithm-trained neural-surrogate methodology, including the finite-element ground-truth model, the surrogate architecture, and the optimization loop. Section 4 describes the case-study configuration and experimental protocol. Section 5 reports and discusses the results, including a discussion of the framework's limitations, and Section 6 concludes the paper and outlines future work.

2. Related Work

2.1 Structural and Topology Optimization of Robotic-Arm Links

Several recent studies have applied finite-element-based topology and shape optimization directly to robotic-arm links in order to reduce mass without violating strength or stiffness requirements. Li and Li [3] combined kinematic and finite-element analysis to redesign a hydraulically-actuated multifunctional robotic arm, reporting a 14.3% mass reduction while retaining strength and stiffness margins. Pang et al. [4] applied variable-density topology optimization to the forearm of a welding robot to suppress actuator-induced vibration while reducing mass, and Ali et al. [5] designed and validated, through finite-element analysis, a low-cost aluminium 5-DOF robotic arm for material-handling applications in small manufacturing settings. Kouritem et al. [6] formulated a multi-objective design problem for an industrial robot arm that jointly selects cross-sectional geometry and material to minimize initial and running cost subject to stress and vibration constraints, using a genetic algorithm to optimize the resulting trajectory-dependent objective. Wu et al. [7] proposed a design-dependent topology-optimization framework for multi-component robotic arms that explicitly accounts for time-varying inertial loading during motion, demonstrating that neglecting motion-dependent loads can misrepresent the true optimal topology. John et al. [8] introduced a carbon-fibre-reinforced-polymer (CFRP) tube-array robot-link concept validated by finite-element analysis, reporting substantial mass reduction with improved dynamic stiffness relative to conventional metallic links, and Hou and Zhang [9] similarly replaced a steel industrial-robot arm with a thermoplastic-CFRP composite boom optimized layer-by-layer under an equal-stiffness criterion. Batista et al. [10] combined topology optimization with additive-manufacturing lattice infill to redesign a robotic-arm link, reporting a 29.3% mass reduction relative to a generic solid design at the cost of a moderate increase in peak equivalent stress. These studies consistently show that lightweight robotic-arm redesign is achievable through iterative structural optimization and/or composite-material substitution, but in every case the optimization loop calls a finite-element solver directly and material selection, where considered at all, is treated qualitatively rather than as a variable optimized jointly with geometry -- both dependencies the present paper replaces with a trained surrogate and a formal discrete decision variable, respectively.

2.2 Genetic and Evolutionary Algorithms in Structural Design

Pastor et al. [11] used a genetic algorithm to compare straight, rounded, and curved link-module geometries for task-based manipulator synthesis, showing that simple, few-parameter modular links can outperform more elaborate ones -- a finding consistent with the compact six-variable parameterization adopted in Section 3. The non-dominated sorting genetic algorithm (NSGA-II) and its successor NSGA-III remain the dominant multi-objective evolutionary engines for constrained structural design problems, and several recent studies apply them to problems structurally similar to the present one: Wu et al. [21] used an improved NSGA-II with a modified crowding-distance operator to optimize an automotive body-in-white structure for minimum mass subject to torsional-stiffness and frequency increments; Zhao et al. [22] applied an improved NSGA-II to the multi-objective optimization of a main bearing assembly structure; Jia et al. [23] combined NSGA-III with generative design to optimize a dual-robot-gripper unloading device under multiple lightweight-design objectives; and Chen et al. [24] coupled NSGA-III with a radial-basis-function surrogate model and entropy-weighted TOPSIS ranking to optimize a body-in-white structure, illustrating -- as does the present paper -- that a trained surrogate can be embedded directly inside the NSGA family rather than calling a finite-element solver at every evaluation. Ma et al. [25] provide a comprehensive, recent survey of the NSGA-II family, its variants, and its engineering applications, while Selvarajan [26] reviews the broader landscape of modern metaheuristic and evolutionary optimization techniques, including genetic algorithms, underlying the constrained NSGA-II implementation adopted in Section 3.8.

2.3 Neural-Network Surrogate Models for Structural Analysis

Data-driven surrogate modelling has become practical for structural-analysis tasks that were previously the sole domain of finite-element solvers, a trend documented in recent deep-learning surveys [27]. Barmada et al. [12] trained a deep-learning surrogate to predict topology-optimization outcomes directly, bypassing the iterative optimality-criteria loop. Ates and Gorguluarslan [13] proposed a two-stage convolutional encoder-decoder architecture that improves the reliability of deep-learning-based topology-optimization predictions relative to single-stage networks. Shao et al. [14] trained fully-connected and recurrent surrogate models on finite-element impact simulations, predicting stress distribution and stress sequence at a fraction of the simulation cost. Hashemi et al. [15] proposed a machine-learning-based surrogate finite element model for the dynamic response of mechanical systems under parametric uncertainty. Most directly relevant to the present study's ground-truth model (Section 3.5), Vurtur Badarinath et al. [17] developed and benchmarked several machine-learning surrogates -- including artificial neural networks, random forests, and gradient-boosted trees -- for the time-varying response of a one-dimensional beam structure, finding neural networks the most accurate; the present paper extends this one-dimensional-beam-surrogate strategy to a stepped, two-module robotic-arm link chain with simultaneous geometry and material selection. Rezapourian et al. [18] trained a multi-output feed-forward neural-network surrogate on thousands of finite-element simulations of architected titanium metamaterials to predict elastic modulus, strength, and energy absorption directly from design parameters, a scale and structure of surrogate very similar to that adopted in Section 3.7. Most directly relevant to the present study's optimization loop, Palmeri and Laurenzi [16] trained neural-network surrogates on finite-element buckling simulations of thin-walled collapsible tubular masts, achieving coefficients of determination between 0.963 and 0.999, and used

the trained surrogate inside NSGA-II to obtain a multi-objective Pareto front of buckling-optimized cross-sections -- a methodology this paper adapts and extends to modular robotic-arm link design with discrete material selection.

2.4 Surrogate-Assisted Evolutionary Optimization

He et al. [19] reviewed the state of the art in surrogate-assisted evolutionary algorithms (SAEAs) for expensive optimization problems, classifying methods by objective and constraint structure and noting that infill/re-evaluation criteria -- periodically checking surrogate predictions against the true simulator -- remain essential for trustworthy convergence, a recommendation this paper follows through its explicit Pareto-front re-validation step (Section 3.9). Liang et al. [20] more recently surveyed SAEA model-construction and model-management strategies, corroborating that the reduction in expensive-evaluation count, rather than raw wall-clock speed, is the primary and most transferable benefit reported across the SAEA literature a conclusion consistent with the computational-cost analysis in Section 5.5.

2.5 Modular and Reconfigurable Manipulator Design and Digital Twins

Wang et al. [1] presented the design and finite-element-based structural verification of a cable-free modular reconfigurable manipulator system, in which interchangeable modules are assembled through a docking mechanism to realize different degrees of freedom for different tasks; finite-element analysis was used to identify and redesign modules that failed to meet strength or stiffness targets. Liang et al. [28] provide a comprehensive recent survey of modular reconfigurable robot mechanisms and design taxonomy, situating module-level structural sizing the problem addressed in this paper -- as a necessary complement to the connector- and actuator-level design questions that dominate that literature. Huang et al. [2] surveyed the broader convergence of artificial intelligence and digital-twin technology in Industry 4.0 manufacturing and robotics, concluding that data-driven models are becoming central to keeping physical assets and their digital representations synchronized and rapidly re-configurable a trend that motivates using a trained, reusable surrogate, rather than a one-off simulation, as the structural-evaluation component of a modular arm's design pipeline. Finally, the Denavit-Hartenberg convention used throughout the robotics literature to describe serial-chain kinematics also underlies the geometric description of the module chain adopted in Section 3; Karhan and Bingül [29] recently developed an automated toolbox for deriving Denavit-Hartenberg parameters directly from a robot's zero configuration, illustrating the convention's continued centrality to serial-manipulator modelling. A broader review of intelligence-algorithm applications across engineering domains is given by Zeng et al. [30].

2.6 Synthesis and Research Gap

Table 1 summarizes how the reviewed studies position themselves along three axes that jointly define the scope of the present paper: (i) whether the structure under study is a robotic-arm link or module, (ii) whether the optimization search is evolutionary, and (iii) whether the fitness/response evaluation is provided by a trained surrogate rather than a direct physics solver. Robotic-arm-specific structural optimization studies [3], [4], [5], [6], [7], [8], [9], [10] populate the region combining arm links with direct FE evaluation; NSGA-family structural-optimization studies [21], [22], [23], [24] populate the region combining generic structures with evolutionary search, occasionally coupled to a surrogate [24];

general surrogate-assisted structural-analysis studies [12], [13], [14], [15], [17], [18] populate the region combining generic structures with surrogate evaluation; and no reviewed study occupies the intersection of all three axes for a modular arm in which material selection is part of the decision vector. This gap motivates the framework proposed in Section 3.

TABLE 1: Positioning of reviewed studies relative to the present work

Study	Structural target	Search method	Response evaluation
Li & Li [3]	Robotic arm (garbage truck)	Parametric FE redesign	Direct FE
Kouritem et al. [6]	Industrial robot arm	Genetic algorithm	Direct FE
Wu et al. [7]	Multi-component robotic arm	Topology optimization (BESO)	Direct FE
Batista et al. [10]	Robot-arm link (AM lattice)	Topology + lattice optimization	Direct FE
Wu et al. [21]	Automotive BIW	Improved NSGA-II	Direct FE
Chen et al. [24]	Automotive BIW	NSGA-III	RBF surrogate
Barmada et al. [12]	Generic structure	Surrogate-driven design	NN surrogate
Vurtur Badarinath et al. [17]	1-D beam structure	Prediction only	NN / RF / GBT surrogate
Palmeri & Laurenzi [16]	Collapsible tubular mast	NSGA-II	NN surrogate
This study	Modular robot-arm link chain	NSGA-II	GA-trained NN surrogate

Note: FE = finite element; AM = additive manufacturing; BIW = body-in-white; BESO = bi-directional evolutionary structural optimization; NN = neural network; RF = random forest; GBT = gradient-boosted trees; RBF = radial basis function.

3. Problem Formulation and Methodology

The design problem addressed in this paper is the simultaneous selection of cross-sectional geometry and material for the two dominant structural modules of a serial robotic-arm link chain: the shoulder-to-elbow module (Link 1) and the elbow-to-wrist module (Link 2), shown schematically in Fig. 2. These two links dominate the arm's structural mass and compliance because the remaining joints and the wrist/tool assembly are comparatively short and are treated as lumped masses, consistent with prior robot-arm optimization studies that found the contribution of the shorter proximal and distal segments to overall stiffness and mass to be small relative to the two principal links [6].

Each module is modelled as a hollow circular tube of outer diameter D and wall thickness t , manufactured in one of three candidate materials (Table 2): aluminium alloy 6061-T6, titanium alloy Ti-6Al-4V, or unidirectional carbon-fibre-reinforced polymer (CFRP-UD). The design vector is therefore

$$x = [D_1, t_1, m_1 D_2, t_2, m_2] \quad (1)$$

where D_i and t_i are the outer diameter and wall thickness of module i ($i = 1, 2$) and m_i is an ordinal-encoded material index, $m_i \in \{0, 1, 2\}$. Bounds on the continuous variables are given in Table 3; the material index is treated as a discrete/categorical variable handled by real-valued relaxation with rounding during optimization, as detailed in Section 3.8.

3.1 Section Geometric Properties

For a hollow circular tube of outer diameter D and wall thickness t , the cross-sectional area A , second moment of area I , and polar moment of area J are

$$A = (\pi/4)[D^2 - (D - 2t)^2] \quad (2)$$

$$I = (\pi/64)[D^4 - (D - 2t)^4] \quad (3)$$

$$J = I_x + I_y = 2I \quad (4)$$

3.2 Objective Functions

The structural mass of the two-module arm chain is

$$M(x) = \rho(m_1)A_1L_1 + \rho(m_2)A_2L_2 + M_j \quad (5)$$

where L_1 and L_2 are the fixed module lengths (Table 3), $\rho(m_i)$ is the density of the material assigned to module i , and $M_j = M_{el} + M_{wr}$ is the combined lumped mass of the elbow-joint actuator/gearbox module and the wrist/tool module. The second objective is the tip deflection $\delta(\mathbf{x})$ of the arm under the worst-case fully-extended horizontal load case described in Section 3.4. The multi-objective design problem is

$$\text{minimize } F(x) = [M(x), \delta(x)] \quad (6)$$

subject to

$$FoS_j(x) = \sigma_y(m_i)/\sigma_{ym_j}(x) \geq FoS_{min}, \quad i = 1, 2 \quad (7)$$

$$f_1(x) \geq f_{1,min} \quad (8)$$

$$\delta(x) \leq \delta_{max} \quad (9)$$

$$t_i/D_i \geq \beta_{min}, \quad i = 1, 2; x \in [x^L, x^U] \quad (10)$$

where $FoS_i(\mathbf{x}) = \sigma_y(m_i)/\sigma_{vm,i}(\mathbf{x})$ is the factor of safety at the critical cross-section of module i , $f_1(\mathbf{x})$ is the first natural frequency of the arm, δ_{max} is the maximum allowable tip deflection, and β_{min} is a minimum wall-thickness-to-diameter ratio introduced to avoid local shell-buckling failure modes that are not otherwise captured by the beam-theory stress calculation. Table 3 lists the numerical values $FoS_{min} = 2.5$, $f_{1,min} = 25$ Hz, $\delta_{max} = 1.5$ mm, and $\beta_{min} = 1/40$ adopted for the case study in Section 4; these values are representative of typical collaborative-robot structural design targets reported in the literature reviewed in Section 2 [6], [7] rather than the specification of any particular commercial product.

TABLE 2: Candidate module materials and properties

Material	E (GPa)	ρ (kg/m ³)	σ_y / allow. (MPa)	Cost index	E/ ρ (MJ/kg)	σ_y/ρ (MJ/kg)
Al 6061-T6	68.9	2700	276	1.0	25.5	0.102
Ti-6Al-4V	113.8	4430	880	4.5	25.7	0.199
CFRP-UD	135.0	1600	600	8.0	84.4	0.375

Note: cost index is a dimensionless relative material-cost factor (aluminium = 1.0); E/ ρ and σ_y/ρ are specific modulus and specific strength.

TABLE 3: Design variables, fixed parameters, and constraint thresholds

Quantity	Symbol	Value / range	Unit	Type
Link 1 outer diameter	D_1	35 – 90	mm	Variable
Link 1 wall thickness	t_1	1.5 – 8.0	mm	Variable
Link 1 material index	m_1	0 (Al) – 2 (CFRP)	–	Variable
Link 2 outer diameter	D_2	25 – 70	mm	Variable
Link 2 wall thickness	t_2	1.2 – 6.0	mm	Variable
Link 2 material index	m_2	0 (Al) – 2 (CFRP)	–	Variable
Link 1 / Link 2 length	L_1, L_2	0.35 / 0.30	m	Fixed
Elbow joint module mass	M_{el}	1.2	kg	Fixed
Wrist / tool module mass	M_{wr}	1.5	kg	Fixed
Rated payload	P_r	5.0	kg	Fixed
Dynamic load factor	DLF	1.5	–	Fixed
Payload eccentricity	e	30	mm	Fixed
Min. factor of safety	FoS_{min}	2.5	–	Constraint
Min. first natural frequency	$f_{1,min}$	25	Hz	Constraint
Max. tip deflection	δ_{max}	1.5	mm	Constraint
Min. thickness/diameter ratio	β_{min}	1/40	–	Constraint

3.3 Overall Framework

Fig. 1 summarizes the proposed genetic-algorithm-trained neural-surrogate framework. Starting from the parametric modular-arm model of Section 3, a Latin-Hypercube design of experiments samples the continuous design variables while the material indices are drawn independently to ensure balanced coverage of all material combinations (Section 3.6). Each sampled design is evaluated by a one-dimensional finite-element beam model that plays the role of the high-fidelity, computationally-expensive simulator in this study (Section 3.5); this generates the labelled dataset used to train the neural-network surrogate (Section 3.7). The trained surrogate is then embedded as the fitness evaluator inside a constrained NSGA-II optimization loop (Section 3.8), and the resulting Pareto-optimal design set is independently re-evaluated with the ground-truth finite-element model to certify feasibility before any design is accepted (Section 3.9). A dashed feedback path in Fig. 1 indicates where an adaptive infill/re-training step could route poorly-predicted candidates back into the training set in a fully online deployment; the present study evaluates the single-shot (offline-trained) variant of the framework and discusses the online extension in Section 5.8.




Fig. 1. Genetic-algorithm-trained neural-surrogate design-optimization framework.

3.4 Parametric Arm Model and Loading

Fig. 2 shows the two-module link chain in the fully-extended horizontal pose, which is the classical worst-case static loading condition for robot-arm structural design because it maximizes the bending-moment arm of both the self-weight and the tip payload [6]. The shoulder joint is treated as a rigid cantilever support; the elbow-joint actuator/gearbox module (mass $M_{el} = 1.2$ kg) is represented as a lumped point mass at the module boundary $x = L_1$; and the wrist/tool module (mass $M_{w_r} = 1.5$ kg) together with the rated payload $P_r = 5$ kg, amplified by a dynamic load factor $DLF = 1.5$ to account for acceleration and impact loading during motion, is applied as a lumped point load at the tip, $x = L_1 + L_2$. A lateral payload eccentricity of $e = 30$ mm is additionally assumed to generate a constant reactive torque $T = P_{tip} \cdot e$ along both modules, producing a secondary torsional-shear stress case that is combined with the bending stress through the von Mises criterion.




Fig. 2. Two-module structural arm chain analysed in the case study (shoulder-to-elbow and elbow-to-wrist modules, fully-extended pose).

3.5 Ground-Truth Structural Evaluation: One-Dimensional Finite-Element Beam Model

Because slender robotic-arm links are well approximated as beams in the structural-optimization literature reviewed in Section 2 [3], [6], [7], the ground-truth structural response used to generate training data and to validate the final Pareto-optimal designs is computed with a two-node Euler–Bernoulli finite-element beam model rather than a full three-dimensional solid finite-element model. This choice keeps the per-evaluation cost low enough to generate a dataset of thousands of samples while retaining the essential physics — distributed self-weight, concentrated joint and payload masses, a stepped cross-section at the module boundary, and free-vibration behaviour that governs the design trade-off of interest. Each module is discretized into 20 two-node beam elements (40 elements, 41 nodes, 82 degrees of freedom in total), and each element contributes the standard Euler–Bernoulli stiffness matrix

$$K_e = (EI/I_e^3) \begin{bmatrix} [12, 6I_e, -12, 6I_e]; [6I_e, 4I_e^2, -6I_e, 2I_e^2]; [-12, -6I_e, 12, -6I_e]; [6I_e, 2I_e^2, -6I_e, 4I_e^2] \end{bmatrix} \quad (11)$$

where E and I are the elastic modulus and second moment of area of the element's parent module and I_e is the element length. Element stiffness matrices are assembled into the global stiffness matrix K , distributed self-weight is converted to statically-equivalent consistent nodal loads, and the elbow and tip point loads are added directly to the corresponding degrees of freedom. After applying the fixed-support boundary condition at the shoulder, the static equilibrium equation

$$Kd = F \quad (12)$$

is solved for the nodal displacement/rotation vector $\mathbf{d}$, from which the tip deflection δ is read directly as the transverse displacement of the final node. Because the arm is a statically determinate cantilever, the bending-moment distribution needed for stress evaluation is obtained directly from statics rather

than from the finite-element solution, avoiding mesh-dependent stress-recovery error; the root moment (governing Link 1) and step moment (governing Link 2, evaluated just past the module boundary) are

$$M_{root} = \Sigma(\text{load}) \cdot (\text{distance from root}) = w_1 L_1 (L_1/2) + w_2 L_2 (L_1 + L_2/2) + F_{e1} L_1 + F_{tip} (L_1 + L_2) \quad (13a)$$

$$M_{step} = w_2 L_2 (L_2/2) + F_{tip} L_2 \quad (13b)$$

where w_i is the self-weight per unit length of module i , $F_{el} = M_{el}g$, and $F_{tip} = (M_{w_r} + P_r)g \cdot \text{DLF}$. The root and step-section von Mises stresses combine bending and torsional shear,

$$\sigma_{vm,i} = \sqrt{(\sigma_{en,i}^b)^2 + 3\tau_j^2}, \sigma_{en,i}^b = M_i (D_i/2)/I_i, \tau_j = T (D_i/2)/J_i \quad (14)$$

and the factor of safety at each section is $\text{FoS}_i = \sigma_y(m_i)/\sigma_{vm,i}$. Finally, the first natural frequency is obtained by solving the generalized eigenvalue problem

$$(K - \omega^2 M_{lump})\phi = 0 \quad (15)$$

where M_{lump} is a lumped mass matrix assembled from each element's distributed mass plus the elbow and tip point masses, and $f_1 = \omega_1/2\pi$ is computed from the lowest non-zero eigenvalue ω_1^2 . Algorithm 1 summarizes the complete ground-truth evaluation procedure.

Algorithm 1 Ground-Truth Structural Evaluation

```

Input: design vector  $x = [D1, t1, m1, D2, t2, m2]$ 
Output:  $M, \text{delta}, \text{sigma\_vm1}, \text{sigma\_vm2}, \text{FoS1}, \text{FoS2}, f1$ 
1: Compute  $A1, I1, J1$  and  $A2, I2, J2$  from  $(D1, t1), (D2, t2)$  [Eqs. 2-4]
2: Assemble global stiffness matrix  $K$  from 40 two-node Euler-Bernoulli beam elements [Eq. 11]
3: Assemble consistent nodal load vector  $F$  from distributed self-weight + elbow point load + tip point load
4: Apply fixed-support boundary condition at shoulder node
5: Solve  $K d = F$  for nodal displacement vector  $d$  [Eq. 12]
6:  $\text{delta} \leftarrow$  transverse displacement at tip node
7: Compute  $M_{root}, M_{step}$  from statics [Eq. 13]
8:  $\text{sigma\_vm1}, \text{sigma\_vm2} \leftarrow$  von Mises stress [Eq. 14]
    $\text{FoS1} \leftarrow \text{sigma}_y(m1)/\text{sigma\_vm1}$  ;  $\text{FoS2} \leftarrow \text{sigma}_y(m2)/\text{sigma\_vm2}$ 
9: Assemble lumped mass matrix  $M_{lump}$ 
10: Solve  $(K - \omega^2 M_{lump}) \phi = 0$  for lowest  $\omega_{\text{eig1}}$  [Eq. 15]
     $f1 \leftarrow \omega_{\text{eig1}} / (2\pi)$ 
11:  $M \leftarrow$  closed-form structural mass [Eq. 5]
12: return  $M, \text{delta}, \text{sigma\_vm1}, \text{sigma\_vm2}, \text{FoS1}, \text{FoS2}, f1$ 

```

3.6 Design of Experiments and Dataset Generation

A space-filling Latin Hypercube Sampling (LHS) design is used to generate 1,200 training-pool designs spanning the continuous variable bounds of Table 3, with material indices m_1 and m_2 drawn independently and uniformly over the three candidate materials to ensure that every material combination is well represented. An independent set of 300 test designs is generated with a different

random seed to provide an unbiased estimate of surrogate generalization performance (Section 5.1). Each of the 1,500 designs is evaluated with the ground-truth model of Section 3.5; feasibility with respect to the constraints of Eqs. (7)–(10) is recorded for reference but is not used to filter the dataset, since the surrogate must also learn to recognize infeasible regions of the design space. The resulting dataset spans structural masses from 2.90 kg to 6.46 kg, tip deflections from 0.06 mm to 4.37 mm, and first natural frequencies from 9.3 Hz to 78.1 Hz, providing broad coverage on both sides of every constraint boundary.

3.7 Neural-Network Surrogate Architecture and Training

A fully-connected feed-forward neural network (multilayer perceptron) is trained to predict the three structurally-expensive responses tip deflection, root and step von Mises stress, and first natural frequency directly from the six-dimensional design vector; structural mass is retained as a closed-form function (Eq. 5) rather than a learned output, since it is inexpensive to evaluate exactly and does not benefit from surrogate approximation. Fig. 3 shows the adopted architecture: an input layer of 6 neurons, three hidden layers of 64, 64, and 32 neurons with rectified-linear-unit (ReLU) activation, and an output layer of 4 neurons with linear activation. For an input x , each hidden layer computes

$$a^{(l)} = \text{RELU}(W^{(l)}a^{(l-1)} + b^{(l)}), \quad a^{(0)} = x \text{ (standardized input)} \quad (16)$$

with $a^{(0)} = x$, and the network output is $\hat{y} = W^{(L)}a^{(L-1)} + b^{(L)}$. Inputs and outputs are standardized (zero mean, unit variance) prior to training using statistics computed on the training split only. The network is trained with the Adam optimizer to minimize the mean-squared error

$$L(\theta) = (1/N) \sum_{n=1}^N \|y_n - \hat{y}_n\|^2 \quad (17)$$

with an L2 weight-decay coefficient of 10^{-4} , an initial learning rate of 10^{-3} , and a batch size of 32. The 1,200-sample training pool is split into 1,000 training and 200 validation designs; training proceeds for up to 400 epochs with early stopping (patience of 25 epochs) on the validation mean-squared error, and the network weights corresponding to the lowest validation error (attained at epoch 100) are retained. Table 4 summarizes the surrogate architecture and training hyperparameters.




Fig. 3. Feed-forward neural-network surrogate architecture (6-64-64-32-4).

TABLE 4: Neural-network surrogate architecture and training hyperparameters

Parameter	Value
Input dimension	6 (D1, t1, m1, D2, t2, m2)
Hidden layers	64, 64, 32 (ReLU activation)
Output dimension	4 (δ , $\sigma_{vm,1}$, $\sigma_{vm,2}$, f_1)
Optimizer	Adam
L2 weight decay (α)	1×10^{-4}
Initial learning rate	1×10^{-3}
Batch size	32
Maximum epochs	400
Early-stopping patience	25 epochs
Best epoch (lowest validation MSE)	100
Training samples	1,000
Validation samples	200
Independent test samples	300

3.8 Genetic Algorithm (NSGA-II) Optimization Loop

The trained surrogate is embedded as the fitness evaluator inside a constrained implementation of NSGA-II [25] developed for this study. For a candidate design x , the surrogate predicts $(\hat{\delta}, \hat{\sigma}_{vm,1}, \hat{\sigma}_{vm,2}, \hat{f}_1) = f_{\theta}(x)$ in normalized space; these are de-normalized, combined with the closed-form mass $M(x)$ of Eq. (5), and used to evaluate the two objectives of Eq. (6) and the constraint set of Eqs. (7)–(10). Constraint satisfaction is aggregated into a single non-negative constraint-violation measure

$$CV(x) = \sum_k \max(0, g_k(x)/|g_{k,ref}|) \quad (18)$$

where each $g_k(x) \leq 0$ is a normalized form of one of Eqs. (7)–(10) and $g_{k,ref}$ is the corresponding constraint bound used purely for scaling; $CV(x) = 0$ denotes a feasible design. Selection uses Deb's constrained-tournament rule: a feasible design always dominates an infeasible one, between two

infeasible designs the one with lower $CV(x)$ wins, and between two feasible designs standard Pareto dominance and crowding distance apply [25]. New candidates are generated from selected parents by simulated binary crossover (SBX),

$$c_1 = \frac{1}{2} [(x_1 + x_2) - \beta q(x_2 - x_1)], \quad c_2 = \frac{1}{2} [(x_1 + x_2) + \beta q(x_2 - x_1)] \quad (19)$$

with distribution index $\eta_c = 15$ and crossover probability $pc = 0.9$, followed by polynomial mutation with distribution index $\eta_m = 20$ and per-variable mutation probability $pm = 1/6$. Algorithm 2 summarizes one generation of the surrogate-assisted optimization loop. The optimizer is run with a population of 80 for 150 generations (12,080 surrogate fitness evaluations), and the resulting rank-0 feasible individuals constitute the surrogate-predicted Pareto front reported in Section 5.3.

Algorithm 2 One Generation of Surrogate-Assisted NSGA-II

Input: parent population P_t (size N), trained surrogate f_{θ}
Output: next-generation population P_{t+1}

- 1: for each x in P_t :
- 2: predict $(\delta_{\hat{t}}, \sigma_{\hat{t}}, \sigma_{\hat{t}2}, f_{\hat{t}}) \leftarrow f_{\theta}(x)$
- 3: compute $M(x)$; evaluate $F(x)$, $CV(x)$ [Eqs. 5, 6, 18]
- 4: $Q_t \leftarrow$ empty offspring population
- 5: while $|Q_t| < N$:
- 6: select parents pa, pb via constrained binary tournament
- 7: $(c_1, c_2) \leftarrow$ SBX_crossover(pa, pb ; $\eta_c=15, pc=0.9$) [Eq. 19]
- 8: $c_1 \leftarrow$ polynomial_mutation(c_1 ; $\eta_m=20, pm=1/6$)
- 9: $c_2 \leftarrow$ polynomial_mutation(c_2 ; $\eta_m=20, pm=1/6$)
- 10: append c_1, c_2 to Q_t
- 11: evaluate all x in Q_t via surrogate (as in lines 1-3)
- 12: $R_t \leftarrow P_t \cup Q_t$
- 13: $(F_1, F_2, \dots) \leftarrow$ constrained_fast_non_dominated_sort(R_t)
- 14: $P_{t+1} \leftarrow$ empty ; $i \leftarrow 1$
- 15: while $|P_{t+1}| + |F_i| \leq N$:
- 16: $P_{t+1} \leftarrow P_{t+1} \cup F_i$; $i \leftarrow i+1$
- 17: fill remaining $N - |P_{t+1}|$ slots from F_i ranked by crowding distance (descending)
- 18: return P_{t+1}

3.9 Surrogate Validation and Sensitivity Analysis

Because a surrogate that is accurate on average may still be locally unreliable near a constraint boundary precisely where an optimizer concentrates its search every design on the surrogate-predicted Pareto front is re-evaluated with the ground-truth model of Section 3.5 after optimization terminates, and only designs that remain feasible under this independent check are reported as certified optimal (Section 5.4), consistent with the infill/re-evaluation practice recommended for surrogate-assisted evolutionary algorithms [19]. In addition, a local sensitivity analysis is performed at the best-compromise (“knee-point”) design, defined as the Pareto-optimal solution minimizing the Euclidean distance to the normalized ideal point in objective space. For each continuous design variable x_i , a central finite difference

$$\partial y / \partial x_i = [y(x + h_i e_i) - y(x - h_i e_i)] / (2h_i) \quad (20)$$

with $h_i = 0.02(x_{i,U} - x_{i,L})$ is evaluated on both the closed-form mass function and the trained surrogate, and the result is scaled by the variable's full range ($x_{i,U} - x_{i,L}$) to obtain a dimensionally consistent contribution measure reported in Fig. 8.

4. Case Study and Experimental Setup

The case-study arm follows the geometry and loading of Section 3.4, with module lengths $L_1 = 0.35$ m and $L_2 = 0.30$ m fixed to represent a small-to-medium collaborative-robot reach of approximately 0.65 m across the two principal links (Fig. 2), consistent with the reach envelopes reported for the modular and industrial manipulators reviewed in Section 2 [1], [6]. The three-candidate module materials aluminium 6061-T6, titanium Ti-6Al-4V, and unidirectional CFRP span roughly a 3.3:1 range in specific stiffness (E/ρ) and a 3.7:1 range in specific strength (σ_y/ρ), so that material selection meaningfully interacts with the geometric design variables rather than being dominated by a single obviously superior choice (Table 2).

All software components the ground-truth finite-element beam model, dataset generation, neural-network training, and the NSGA-II optimizer were implemented in Python 3.12 using NumPy 2.4, SciPy 1.17 (Latin Hypercube Sampling and generalized eigenvalue solution), and scikit-learn 1.8 (neural-network training and data standardization); the NSGA-II algorithm itself was implemented from scratch to give full control over the constrained-dominance rule of Section 3.8. All experiments were executed on a single workstation CPU core; no GPU acceleration was required given the small size of the surrogate network. Table 5 lists the complete set of NSGA-II hyperparameters used in both the surrogate-assisted run and the direct FE-in-the-loop control run, which are otherwise identical (same population, generation count, operators, and random seed for population initialization) so that any difference in outcome can be attributed to the fitness-evaluation method rather than to search-algorithm configuration.

Two optimization runs were executed for comparison: (i) the proposed surrogate-assisted NSGA-II, in which every fitness evaluation calls the trained neural-network surrogate of Section 3.7, and (ii) a direct FE-in-the-loop NSGA-II, in which every fitness evaluation instead calls the ground-truth finite-element model of Section 3.5. Both runs share identical population size, generation count, genetic operators, and random seed, isolating the effect of the fitness-evaluation strategy. Following optimization, every design on the surrogate-assisted Pareto front is independently re-evaluated with the ground-truth model to quantify surrogate reliability at the optimum (Section 5.4), and the two runs' convergence histories and computational cost are compared directly (Sections 5.2 and 5.5).

TABLE 5: NSGA-II Optimizer hyperparameters

Parameter	Value
Population size	80
Generations	150
Crossover operator	Simulated binary crossover (SBX)
SBX distribution index (η_c)	15
Crossover probability (p_c)	0.9
Mutation operator	Polynomial mutation
Mutation distribution index (η_m)	20
Mutation probability (p_m)	1/6 (≈ 0.167)

Random seed	3
Total fitness evaluations	12,080

5. Results and Discussion

5.1 Surrogate Model Accuracy

Table 6 reports the test-set accuracy of the trained neural-network surrogate for all four predicted responses, and Fig. 4 shows the corresponding training and validation loss curves. The surrogate attains $R^2 \geq 0.984$ on every response, with the lowest accuracy on tip deflection ($R^2 = 0.984$, MAPE = 6.1%) and the highest on first natural frequency ($R^2 = 0.997$, MAPE = 2.1%). Fig. 5 shows the corresponding parity plots, confirming that predictions track the 1:1 line closely across the full range of the test set with no systematic bias, though the deflection predictions show mildly increased scatter at the largest deflection values, where the underlying mechanical response becomes most sensitive to the thinnest, most flexible geometries in the design space.

TABLE 6: Neural-surrogate accuracy on the independent test set (N = 300)

Response	R^2	RMSE	MAE	MAPE (%)
Tip deflection, δ (mm)	0.984	0.071	0.035	6.1
Root stress, $\sigma_{vm,1}$ (MPa)	0.990	0.687	0.390	4.7
Step stress, $\sigma_{vm,2}$ (MPa)	0.994	0.595	0.384	5.5
First natural frequency, f_1 (Hz)	0.997	0.707	0.548	2.1




Fig. 4. Neural-surrogate training and validation loss.




Fig. 5. Parity between neural-surrogate predictions and ground-truth FE results (held-out test set, $n = 300$).

5.2 Convergence of Surrogate-Assisted vs. Direct FE-in-the-Loop NSGA-II

Fig. 6 compares the best feasible structural mass found by generation for the surrogate-assisted and direct FE-in-the-loop NSGA-II runs. Both runs follow closely overlapping convergence trajectories and terminate within 0.07% of one another (2.945 kg vs. 2.947 kg after 150 generations), indicating that the trained surrogate provides a sufficiently faithful fitness landscape for the genetic algorithm to locate essentially the same optimum as it would find by calling the ground-truth model directly. The feasible fraction of the population reached 100% by the first generation in both runs, reflecting the relatively wide feasible region of the design space established in Section 3.6 (approximately 59% of randomly sampled designs satisfy all constraints).




Fig. 6. Convergence of surrogate-assisted vs. direct FE-in-the-loop NSGA-II (best feasible mass per generation).

5.3 Pareto-Optimal Trade-Offs Between Mass and Tip Deflection

Fig. 7 shows the final Pareto front in the mass-deflection objective plane, comparing the surrogate-predicted front to the same designs re-evaluated with the ground-truth model. The two fronts are in close agreement over the well-populated central region of the trade-off curve and diverge modestly at the stiffest (lowest-deflection) extreme, where the surrogate under-predicts deflection by up to 0.12 mm in absolute terms consistent with the reduced density of very-low-deflection training samples noted in Section 3.6. Table 7 lists five representative designs spanning the Pareto front, labelled A (lightest feasible) through E (stiffest), together with their ground-truth-verified performance. Every candidate on the front selects CFRP-UD for both modules; this outcome is consistent with the material's specific-stiffness and specific-strength advantage summarized in Table 2 and echoes the composite-tube lightweighting trend reported in the robotic-arm literature reviewed in Section 2 [3]–[10]. Relative to a representative mid-range aluminium baseline (D1 = 60 mm, t1 = 4 mm, D2 = 45 mm, t2 = 3 mm; mass 3.686 kg, deflection 0.591 mm, f1 = 25.4 Hz), the lightest ground-truth-feasible Pareto design (Design A) reduces structural mass by 19.9% while satisfying all constraints with a comfortable margin (FoS1 = 39.0, FoS2 = 40.8, f1 = 25.2 Hz), and the knee-point compromise design (Design C) reduces mass by 9.8% while providing substantially larger stiffness and frequency margins (f1 = 48.0 Hz) that would tolerate additional payload or duty-cycle uncertainty not explicitly modelled in the present formulation.

TABLE 7: Representative pareto-optimal designs (GROUND-TRUTH-VERIFIED)

Design	D1 (mm)	t1 (mm)	D2 (mm)	t2 (mm)	Mat.	Mass (kg)	δ (mm)	f1 (Hz)	Δ Mass vs. baseline
A (lightest)	61.8	1.57	47.0	1.23	CFRP	2.951	0.598	25.2	−19.9%
B	72.4	1.86	53.5	1.36	CFRP	3.038	0.325	34.2	−17.6%

C (knee)	90.0	2.53	42.4	4.08	CFRP	3.324	0.165	48.0	−9.8%
D	90.0	5.28	60.8	3.88	CFRP	3.820	0.073	72.3	+3.7%
E (stiffest)	90.0	7.21	70.0	6.00	CFRP	4.329	0.050	87.1	+17.4%

Note: material shown is identical (CFRP-UD) for Link 1 and Link 2 for every listed design; $FoS1$ and $FoS2$ exceed 38 for all five designs (see Section 5.3).




Fig. 7. Pareto-optimal trade-off between structural mass and tip deflection.

5.4 Validation Against Ground-Truth Finite-Element Results

Re-evaluating all 80 designs on the surrogate-assisted Pareto front with the ground-truth model shows that 79 of 80 (98.8%) remain feasible; the single infeasible candidate is the extreme lightest-mass point, whose ground-truth natural frequency (24.63 Hz) falls marginally below the 25 Hz constraint despite being predicted feasible by the surrogate a reminder that the very edge of the feasible region, where training data is naturally sparsest, is exactly where surrogate error is most consequential and where the independent re-validation step of Section 3.9 is indispensable. Because structural mass is computed in closed form rather than learned, the mass values on the Pareto front are exact by construction; the mean absolute error in surrogate-predicted tip deflection across the Pareto set is 0.067 mm, of the same order as the 0.035 mm mean absolute error measured on the independent test set (Table 6), though expressed as a relative percentage this error appears far larger for the stiffest designs simply because their ground-truth deflection approaches zero an artefact of the percentage-error metric rather than a meaningful loss of accuracy, and the reason Table 8 reports absolute rather than relative deflection error.

TABLE 8: Validation of the surrogate-assisted pareto front against ground-truth Fe results

Metric	Value
Pareto-front size	80 designs
Ground-truth-feasible fraction	98.8% (79 / 80)
Mean absolute mass error	0.00% (exact, closed-form)

Mean absolute deflection error	0.067 mm
Maximum absolute deflection error	0.122 mm

5.5 Computational Cost Comparison

Table 9 compares the computational cost of the two optimization strategies. Building and validating the neural-network surrogate consumed 1,500 ground-truth finite-element evaluations (1,200 for the training pool, 300 for the independent test set); a further 80 ground-truth evaluations were used to re-validate the final Pareto front, for a total of 1,580 true structural evaluations. The direct FE-in-the-loop NSGA-II, by contrast, consumed 12,080 true structural evaluations to reach an equivalent optimum — a 7.6-fold reduction in high-fidelity evaluations for the surrogate-assisted pipeline. Because the 1,200-sample training set can be reused without modification for any subsequent re-optimization (for example, under a different constraint set or objective weighting), the amortized reduction in structural-evaluation cost across repeated design studies approaches 10.1-fold (12,080/1,200). It should be emphasized that the wall-clock-time comparison reported in Table 9 (21.5 s vs. 31.3 s) reflects the low per-call cost of the reduced-order one-dimensional beam model adopted in this study and does not represent the wall-clock savings that would be obtained with a full three-dimensional commercial finite-element solver, where a single high-fidelity evaluation can require minutes to hours rather than milliseconds; the count of high-fidelity evaluations avoided, rather than the wall-clock ratio measured here, is therefore the more transferable indicator of the framework's practical benefit and is discussed further in Section 5.8.

TABLE 9: Computational cost comparison

Metric	Surrogate-assisted NSGA-II	Direct FE-in-the-loop NSGA-II
Ground-truth FE evals. (dataset)	1,500	0
Ground-truth FE evals. (validation)	80	12,080
Total ground-truth FE evaluations	1,580	12,080
Surrogate (NN) fitness evaluations	12,080	—
Reduction in ground-truth evaluations	7.6×	1× (reference)
Amortized reduction (reused surrogate)	10.1×	—
Measured wall-clock time	21.5 s	31.3 s
Final best feasible mass	2.945 kg	2.947 kg

5.6 Sensitivity Analysis

Fig. 8 shows the local sensitivity of structural mass and first natural frequency to the four continuous design variables at the knee-point design (Design C). Link 1 wall thickness (t_1) is the dominant driver of structural mass, consistent with its direct proportionality to cross-sectional area at the largest-diameter, highest-self-weight module, while Link 2 outer diameter (D_2) is the dominant driver of the first natural frequency, reflecting the fourth-power dependence of second moment of area on diameter (Eq. 3) acting on the more flexible distal module. Link 1 outer diameter and Link 2 wall thickness have comparatively modest and roughly comparable influence on both responses. Because material selection is a discrete decision rather than a continuously differentiable variable, its influence is instead evident directly in the Pareto-front composition of Section 5.3, where CFRP-UD is selected unanimously.




Fig. 8. Local sensitivity of mass and f1 at the knee-point design.

5.7 Comparison with Related Literature

Table 10 situates the mass reduction obtained in this study against reduction percentages explicitly reported in two of the robotic-arm structural-optimization studies reviewed in Section 2. The 19.9% mass reduction achieved by the surrogate-assisted framework (Design A) exceeds the 14.3% reduction reported for a finite-element-based structural redesign of a multifunctional robotic arm [3], reflecting in part this study's inclusion of material selection alongside geometric resizing, and is comparable to, though somewhat below, the 29.3% reduction reported for a topology-and-lattice-optimized additively-manufactured arm [10], whose design freedom (arbitrary internal lattice topology) exceeds that of the parametric tubular-section representation adopted here. This comparison should be read as indicative rather than a controlled benchmark, since the studies differ in baseline geometry, load case, and manufacturing constraints; its purpose is to confirm that the mass-reduction magnitude obtained by the proposed surrogate-assisted framework is of the same order as and mechanistically consistent with reductions achieved by more computationally expensive direct-FEA optimization studies in the literature.

TABLE 10: Indicative comparison with reported mass reductions in related robotic-arm structural-optimization studies

Study	Structure	Method	Reported mass reduction
Li & Li [3]	Multifunctional robotic arm	FE + structural topology redesign	14.3%
Batista et al. [10]	Robotic-arm link (AM)	Topology + lattice optimization	29.3%
This study (Design A)	Modular 2-link arm chain	GA-trained NN surrogate + NSGA-II	19.9%
This study (Design C)	Modular 2-link arm chain	GA-trained NN surrogate + NSGA-II	9.8% (+92% f1 margin)

5.8 Limitations

The ground-truth model uses a one-dimensional Euler–Bernoulli beam representation, which captures bending, torsion, self-weight, and modal behaviour well for slender tubular links; extending the surrogate to full three-dimensional finite-element data would additionally resolve local stress concentrations at joints and fasteners, and is a natural next step before structural certification. The optimization focuses on the two dominant structural modules, with the remaining joints and tool assembly treated as fixed lumped masses, consistent with common practice in the literature [6]; a fully modular design would extend this treatment to every module. The computational-cost comparison in Section 5.5 reflects the low per-evaluation cost of this reduced-order model, so the practical benefit of the proposed framework scales with the cost of whichever high-fidelity solver it replaces. The surrogate was trained once, offline; its accuracy eases slightly at the sparsely-sampled edge of the feasible region (Section 5.4), and the adaptive infill strategy indicated in Fig. 1 offers a promising way to strengthen reliability there further. Finally, the material and load-case assumptions were chosen to be representative of typical robotic-arm practice; applying the framework to a specific platform would involve substituting the appropriate material catalogue, load cases, and constraint thresholds.

6. Conclusion and Future Work

This paper presented a genetic-algorithm-trained neural-surrogate framework for the structural design optimization of modular robotic-arm links, in which a feed-forward neural network trained on a modest, Latin-Hypercube-sampled finite-element dataset replaces the great majority of expensive structural evaluations inside a constrained NSGA-II optimization loop. Applied to a two-module shoulder-to-elbow and elbow-to-wrist link chain with simultaneous geometry and material selection, the trained surrogate achieved coefficients of determination between 0.984 and 0.997 on an independent test set, and the surrogate-assisted optimizer converged to essentially the same Pareto-optimal mass-deflection trade-off as a matched direct finite-element-in-the-loop genetic algorithm while consuming only a fraction of the high-fidelity structural evaluations. Independent re-validation confirmed that 98.8% of the surrogate-optimized Pareto front remained feasible under the ground-truth model, and the best surrogate-optimized design reduced structural mass by 19.9% relative to a representative baseline while satisfying all stress, frequency, and buckling constraints a reduction consistent in magnitude with prior, more computationally expensive direct-FEA robotic-arm redesign studies.

Future work should extend the ground-truth model from the one-dimensional beam representation adopted here to a full three-dimensional shell or solid finite-element model capable of resolving local stress concentrations at module docking interfaces, retrain the surrogate accordingly, and evaluate the framework on a complete modular arm in which every module including joint housings and the wrist/tool assembly is simultaneously optimized rather than treated as a fixed lumped mass. The adaptive infill strategy indicated in Fig. 1, in which surrogate predictions near the evolving Pareto front are selectively re-evaluated with the ground-truth model and used to periodically re-train the network, is a promising route to further improving surrogate reliability at the constraint boundary without a proportional increase in dataset size, and should be compared quantitatively against the offline-trained variant studied here. Finally, extending the design vector to include dynamic trajectory and control parameters alongside structural geometry would allow the framework to jointly optimize a modular

arm's structural and motion-planning performance within a single genetic-algorithm-trained surrogate pipeline.

Declarations

Ethics approval and consent to participate

Not applicable. This study is a purely computational design-optimization investigation that does not involve human participants, human tissue, animals, or clinical data.

Consent for publication

Not applicable.

Funding

This research received no external funding. The authors conducted the work using institutional computational resources and open-source software libraries (NumPy, SciPy, and scikit-learn).

Competing interests

The authors declare that they have no competing interests.

Availability of data and materials

The Latin-Hypercube design-of-experiments dataset, trained neural-network weights, and all Python source code used to generate the results are available from the corresponding author upon reasonable request.

Authors' contributions

T.O.A. conceived the study, developed the methodology, implemented the computational framework, performed the analyses, and wrote the manuscript. The author read and approved the final manuscript.

Acknowledgment

The authors gratefully acknowledge the use of open-source scientific computing libraries (NumPy, SciPy, and scikit-learn) in the implementation of the computational framework described in this paper.

References

- [1] Y. Wang, J. Li, K. Wang, and S. Wang, "Design and analysis of modular reconfigurable manipulator system," *Mathematics*, vol. 13, no. 7, art. 1103, Mar. 2025, doi: 10.3390/math13071103.
- [2] Z. Huang, Y. Shen, J. Li, M. Fey, and C. Brecher, "A survey on AI-driven digital twins in Industry 4.0: Smart manufacturing and advanced robotics," *Sensors*, vol. 21, no. 19, art. 6340, Sep. 2021, doi: 10.3390/s21196340.
- [3] H. Li and Y. Li, "Finite element analysis and structural optimization design of multifunctional robotic arm for garbage truck," *Front. Mech. Eng.*, vol. 11, art. 1543967, Apr. 2025, doi: 10.3389/fmech.2025.1543967.
- [4] H. Pang, Z. Sun, J. Hu, and F. Yang, "Light-weighting and comparative simulation analysis of the forearm of welding robots," *Actuators*, vol. 13, no. 6, art. 209, Jun. 2024, doi: 10.3390/act13060209.

- [5] Z. Ali, M. F. Sheikh, A. Al Rashid, Z. U. Arif, M. Y. Khalid, R. Umer, and M. Koç, “Design and development of a low-cost 5-DOF robotic arm for lightweight material handling and sorting applications: A case study for small manufacturing industries of Pakistan,” *Results Eng.*, vol. 19, art. 101315, Sep. 2023, doi: 10.1016/j.rineng.2023.101315.
- [6] S. A. Kouritem, M. I. Abouheaf, N. Nahas, and M. Hassan, “A multi-objective optimization design of industrial robot arms,” *Alexandria Eng. J.*, vol. 61, no. 12, pp. 12847–12867, Dec. 2022, doi: 10.1016/j.aej.2022.06.052.
- [7] C. Wu, Y. Xu, J. Fang, G. Sun, G. P. Steven, and Q. Li, “Topology optimization for multi-component robotic arms under time-varying loads,” *Struct. Multidiscip. Optim.*, Sep. 2025, doi: 10.1007/s00158-025-04129-1.
- [8] K. John, S. Fletcher, T. Furness, A. Longstaff, and P. Needham, “Novel robot arm concept for lightweighting high-performance machinery,” *J. Mach. Des. Autom. Intell.*, vol. 2, no. 1, art. 20230006, 2024, doi: 10.1515/jmdai-2023-0006.
- [9] H. Hou and P. Zhang, “Lightweight design of industrial robot arm based on thermoplastic carbon fiber reinforced materials,” *J. Phys.: Conf. Ser.*, vol. 2954, art. 012038, 2025, doi: 10.1088/1742-6596/2954/1/012038.
- [10] R. C. Batista, A. Agarwal, A. Gurung, A. Kumar, F. Altarazi, N. Dogra, H. M. Vishwanatha, D. S. Chiniwar, and A. Agrawal, “Topological and lattice-based AM optimization for improving the structural efficiency of robotic arms,” *Front. Mech. Eng.*, vol. 10, art. 1422539, Jun. 2024, doi: 10.3389/fmech.2024.1422539.
- [11] R. Pastor, Z. Bobovský, D. Huczala, and S. Grushko, “Genetic optimization of a manipulator: Comparison between straight, rounded, and curved mechanism links,” *Appl. Sci.*, vol. 11, no. 6, art. 2471, Mar. 2021, doi: 10.3390/app11062471.
- [12] S. Barmada, N. Fontana, A. Formisano, D. Thomopulos, and M. Tucci, “A deep learning surrogate model for topology optimization,” *IEEE Trans. Magn.*, vol. 57, no. 6, pp. 1–4, Jun. 2021, doi: 10.1109/TMAG.2021.3063470.
- [13] G. C. Ates and R. M. Gorguluarslan, “Two-stage convolutional encoder–decoder network to improve the performance and reliability of deep learning models for topology optimization,” *Struct. Multidiscip. Optim.*, vol. 63, no. 4, pp. 1927–1950, Apr. 2021, doi: 10.1007/s00158-020-02788-w.
- [14] M. Shao, J. Chen, T. Wang, F. Tao, J. Du, D. Zhang, X. Wang, and X. Tang, “Accelerating analysis for structure design via deep learning surrogate models,” *Adv. Intell. Syst.*, vol. 5, no. 2, art. 2200099, Feb. 2023, doi: 10.1002/aisy.202200099.
- [15] A. Hashemi, J. Jang, and J. Beheshti, “A machine learning-based surrogate finite element model for estimating dynamic response of mechanical systems,” *IEEE Access*, vol. 11, pp. 54509–54525, 2023, doi: 10.1109/ACCESS.2023.3282453.
- [16] F. Palmeri and S. Laurenzi, “Neural network-based surrogate modeling for buckling performance optimization of lightweight-composite collapsible tubular masts,” *Biomimetics*, vol. 9, no. 8, art. 494, Aug. 2024, doi: 10.3390/biomimetics9080494.

- [17] P. Vurtur Badarinath, M. Chierichetti, and F. Davoudi Kakhki, "A machine learning approach as a surrogate for a finite element analysis: Status of research and application to one dimensional systems," *Sensors*, vol. 21, no. 5, art. 1654, Feb. 2021, doi: 10.3390/s21051654.
- [18] M. Rezapourian, A. Cheloe Darabi, M. Khoshbin, and I. Hussainova, "Surrogate-model prediction of mechanical response in architected Ti6Al4V cylindrical TPMS metamaterials," *Metals*, vol. 15, no. 12, art. 1372, Dec. 2025, doi: 10.3390/met15121372.
- [19] C. He, Y. Zhang, D. Gong, and X. Ji, "A review of surrogate-assisted evolutionary algorithms for expensive optimization problems," *Expert Syst. Appl.*, vol. 217, art. 119495, May 2023, doi: 10.1016/j.eswa.2022.119495.
- [20] J. Liang, Y. Lou, M. Yu, Y. Bi, and K. Yu, "A survey of surrogate-assisted evolutionary algorithms for expensive optimization," *J. Membr. Comput.*, vol. 7, pp. 108–127, 2025, doi: 10.1007/s41965-024-00165-w.
- [21] X. Wu, Z. Chen, T. Liu, H. Song, Z. Wang, and W. Shi, "Improved NSGA-II and its application in BIW structure optimization," *Adv. Mech. Eng.*, vol. 15, no. 2, Feb. 2023, doi: 10.1177/16878132221150588.
- [22] X. Zhao, T. Su, X. Liu, and X. Zhang, "Multi-objective optimization of main bearing assembly structure based on improved NSGA-II," *Energy Sci. Eng.*, vol. 10, no. 1, pp. 43–63, Jan. 2022, doi: 10.1002/ese3.1003.
- [23] J. Jia, X. Sun, T. Liu, J. Tang, J. Wang, and X. Hu, "Structural optimization design of dual robot gripper unloading device based on intelligent optimization algorithms and generative design," *Sensors*, vol. 23, no. 19, art. 8298, Oct. 2023, doi: 10.3390/s23198298.
- [24] H. Chen, C. Lu, L. Feng, Z. Liu, Y. Sun, and W. Chen, "Structural optimization design of BIW using NSGA-III and entropy weighted TOPSIS methods," *Adv. Mech. Eng.*, vol. 15, no. 12, Dec. 2023, doi: 10.1177/16878132231220351.
- [25] H. Ma, Y. Zhang, S. Sun, T. Liu, and Y. Shan, "A comprehensive survey on NSGA-II for multi-objective optimization and applications," *Artif. Intell. Rev.*, vol. 56, no. 12, pp. 15217–15270, 2023, doi: 10.1007/s10462-023-10526-z.
- [26] S. Selvarajan, "A comprehensive study on modern optimization techniques for engineering applications," *Artif. Intell. Rev.*, vol. 57, pp. 1–52, 2024, doi: 10.1007/s10462-024-10829-9.
- [27] M. H. Mohd Noor and A. O. Ige, "A survey on state-of-the-art deep learning applications and challenges," *Eng. Appl. Artif. Intell.*, vol. 159, part B, art. 111225, Nov. 2025, doi: 10.1016/j.engappai.2025.111225.
- [28] G. Liang, D. Wu, Y. Tu, and T. L. Lam, "Decoding modular reconfigurable robots: A survey on mechanisms and design," *Int. J. Robot. Res.*, 2025, doi: 10.1177/02783649241283847.
- [29] H. Karhan and Z. Bingül, "Automatic determination of the Denavit–Hartenberg parameters for the forward kinematics of all serial robots: Novel kinematics toolbox," *Machines*, vol. 13, no. 10, art. 944, Oct. 2025, doi: 10.3390/machines13100944.
- [30] Y. Zeng, B. Liao, Z. Li, C. Hua, and S. Li, "A comprehensive review of recent advances on intelligence algorithms and information engineering applications," *IEEE Access*, vol. 12, pp. 135886–135912, 2024, doi: 10.1109/ACCESS.2024.3461756.