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.

Fidelity-Aware Inverse Dynamics of a Two-Section Tendon-Driven Continuum Robot: Cosserat-Rod Modeling, RK4/Newmark Cross-Verification, and Eleven-Metaheuristic PID Tuning

Somia Guechi
University of M'Sila
Selman Djeffal
djeffalselman@gmail.com

National Polytechnic School of Constantine (ENPC)

Research Article

Keywords: Continuum robot, Cosserat rod, inverse dynamics, tendon-driven robot, PID tuning, metaheuristic optimization, Runge–Kutta, Newmark-β, model fidelity

Posted Date: October 15th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11009949/v1

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

Additional Declarations: No competing interests reported.

Fidelity-Aware Inverse Dynamics of a Two-Section Tendon-Driven Continuum Robot: Cosserat-Rod Modeling, RK4/Newmark Cross-Verification, and Eleven-Metaheuristic PID Tuning

Somia Guechi [1] and Selman Djeffal [2]

[1]Department of Mathematics, Faculty of Mathematics and Computer Science, University of M'Sila, M'Sila, Algeria

[2]Department of Mechanical Engineering, National Polytechnic School of Constantine (ENPC), Constantine, Algeria

Corresponding author: Selman Djeffal (djeffalselman@gmail.com)

Abstract

Tendon-driven continuum robots offer large shape adaptability but remain difficult to control because distributed elasticity, configuration-dependent inertia, tendon constraints, numerical integration, and controller tuning interact in a strongly nonlinear manner. This paper develops a fidelity-aware computational framework for the inverse dynamics and closed-loop control of a two-section tendon-driven continuum robot represented by a reduced shear-inextensible Cosserat rod. The generalized curvature state contains four bending coordinates, while the actuation layer maps six nonnegative tendon tensions to generalized bending effort under preload and saturation constraints. A model-based inverse-dynamics term is combined with section-wise PID feedback, and the resulting six gains are tuned

independently by eleven established population-based optimizers: particle swarm optimization (PSO), genetic algorithm (GA), artificial bee colony (ABC), ant colony optimization (ACO), butterfly optimization algorithm (BOA), whale optimization algorithm (WOA), grey wolf optimizer (GWO), teaching-learning-based optimization (TLBO), Harris hawks optimization (HHO), pelican optimization algorithm (POA), and cuckoo search (CS). To separate optimization effects from numerical artifacts, the same tuned controllers are evaluated with fourth-order Runge–Kutta (RK4) and implicit Newmark-β integration. Four disturbance cases are additionally considered: nominal motion, generalized payload loading, measurement noise, and a combined payload–impulse–noise condition. The reference-linearized benchmark identifies GWO and POA as the strongest RK4 controllers, with nominal tip root-mean-square error (RMSE) values of 24.699 and 24.945 mm, respectively, compared with approximately 155–157 mm for several weakly tuned alternatives. Under the combined disturbance, POA and GWO retain tip RMSE values of 25.062 and 27.265 mm. However, replay of the GWO controller with fully configuration-dependent Cosserat mass, gravity, and Coriolis terms increases the tip RMSE to 400.589 mm. This large fidelity gap is a central result rather than a hidden numerical inconvenience: it demonstrates that low optimization cost and good performance on a reference-linearized plant do not, by themselves, establish robust nonlinear control. The framework therefore provides a reproducible benchmark for simultaneously auditing optimizer choice, integrator sensitivity, tendon feasibility, robustness, and model fidelity.

Keywords: Continuum robot; Cosserat rod; inverse dynamics; tendon-driven robot; PID tuning; metaheuristic optimization; Runge–Kutta; Newmark-β; model fidelity

1 Introduction

Continuum robots replace a finite sequence of rigid revolute or prismatic joints with continuously deformable structures that can negotiate confined, uncertain, or contact-rich environments. This geometric advantage is accompanied by a difficult control problem: the same compliance that enables safe and dexterous interaction also couples shape, actuation, inertia, gravity, constitutive behavior, tendon routing, and numerical approximation. The present study addresses this coupled problem from a deliberately computational perspective. Rather than proposing another isolated optimizer or another isolated forward model, it asks a more demanding question: *when an inverse-dynamics PID controller appears successful for a tendon-driven Cosserat robot, does that conclusion remain credible across optimizer choice, time integrator, disturbance condition, tendon constraints, and model fidelity?*

The recent continuum-robot literature provides a strong basis for this question. Russo et al. [1] organized the broad design, modeling, sensing, actuation, and control landscape and showed how the field has evolved from constant-curvature geometric descriptions toward more general mechanics-based and data-assisted formulations. Extending the design perspective, Alandoli et al. [2] reviewed extensible architectures and highlighted the close relationship between actuation mechanism, variable stiffness, and attainable control authority. Research by Fan et al. [3] then examined stiffening strategies in detail, emphasizing that controllability cannot be separated from the mechanics of deformation. On the sensing side, Sincak et al. [4] synthesized proprioceptive and exteroceptive sensing technologies and reinforced the need to distinguish true plant behavior from errors introduced by state estimation and measurement uncertainty.

Optimization has also become increasingly relevant because continuum-robot inverse problems and controller parameters are often nonconvex and poorly conditioned. Djeffal and Mahfoudi [5] developed a particle-swarm-based inverse-kinematic approach for multi-section continuum robots and compared it with several metaheuristics, illustrating that search strategy can materially change solution quality. In a subsequent design-oriented study, Djeffal et al. [6] linked cable-driven continuum-robot optimization, inverse-kinematic learning, and additive-manufacturing-oriented geometry, demonstrating the value of treating geometry and computation

as a coupled design problem. Moving from inverse kinematics to closed-loop learning, Djeffal et al. [7] used deep deterministic policy-gradient control for a spatial three-section continuum robot. The dynamic side was addressed by Djeffal, Mahfoudi, and Ghoul [8], who solved nonlinear continuum-robot equations with Runge–Kutta integration, a particularly relevant point because numerical integration is not merely an implementation detail when flexible dynamics are strongly nonlinear.

Metaheuristic controller tuning has been explored more directly by Aner et al. [9], who evaluated PSO-PID and PSO-FLC strategies for a continuum-robot model. Analytical control alternatives continue to progress in parallel: Ba et al. [10] developed a differential-kinematics-based control method for tendon-driven continuum robots, while He et al. [11] formulated a model-less visual-control scheme using recurrent-neural-network-based neurodynamic optimization. These works illustrate two complementary directions: exploiting a physical model when sufficiently reliable, or bypassing parts of that model through learning or visual feedback when uncertainty dominates.

Recent studies increasingly confront the nonlinear dynamics directly. Feizi et al. [12] introduced a deep Koopman framework for nonlinear tendon-driven continuum-robot dynamics and control. Alora et al. [13] pursued dominant nonlinear dynamics discovery, showing that reduced representations can be effective when they preserve the behavior relevant to feedback. In contrast, Danesh and Janabi-Sharifi [14] adopted a Cosserat-rod description for backstepping control under large deflections, retaining a stronger connection to continuum mechanics. Fault tolerance and uncertainty have simultaneously moved to the foreground: Liu et al. [15] considered adaptive neural-network fault-tolerant control under actuator faults and external disturbances, whereas Maghooli et al. [16] investigated learning-based control of tendon-driven continuum arms. These contributions make clear that the core challenge is no longer simply reaching a desired shape; it is maintaining credible closed-loop behavior when the assumed model and the realized system diverge.

Mechanical architecture further modifies this challenge. Huertas Nino et al. [17] developed a hybrid tendon-driven architecture designed to avoid torsion under external loading. Yang et al. [18] studied a continuum robot with switchable stiffness, illustrating how physical reconfigura-

tion can change the control landscape itself. From an industrial planning viewpoint, Ouyoucef et al. [19] incorporated durability into trajectory planning and quasi-static control, thereby broadening the performance objective beyond geometric error alone. Obstacle-aware inverse kinematics with experimental validation was reported by Ghoul et al. [20], where teaching–learning-based optimization was used for a variable-length continuum robot. Together, these works motivate performance measures that include actuation burden, feasibility, robustness, and not only endpoint accuracy.

The newest literature is moving toward hybrid optimization, predictive control, and stronger fault-aware formulations. Jabari et al. [21] combined supervised learning and constrained optimization for task-space trajectory generation. Arefinia et al. [22] developed an error-state model-predictive path-integral controller based on Cosserat dynamics with strain parametrization, demonstrating the value of embedding continuum mechanics directly in predictive feedback. In another direction, Jabari et al. [23] proposed prescribed-time fault-tolerant optimal control with composite deep learning for continuum surgical robots. Finally, Djeffal et al. [24] combined machine learning with experimental validation for precise continuum-robot motion, reinforcing the broader trend toward controllers whose credibility is assessed against more than one representation of the plant.

Despite this progress, three methodological gaps remain important. First, many papers compare control laws without explicitly checking whether the ranking changes with the numerical integrator. A controller that is apparently excellent under one integration scheme may exploit numerical damping, discretization, or convergence behavior that is absent under another. Second, optimization studies commonly report the lowest objective value as though it were interchangeable with physical tracking quality, even though the objective may include effort, saturation, transient error, and normalization terms whose trade-offs can lead to very different closed-loop responses. Third, reduced or reference-linearized models are attractive during expensive metaheuristic tuning, but the resulting controller is not automatically reliable under the full configuration-dependent dynamics. The difference between *optimization-model success* and *full-model success* is rarely exposed as a quantitative result.

The present paper addresses these gaps through an integrated benchmark for a two-section,

six-tendon continuum robot. Its main contribution is not a newly named optimizer. Instead, the novelty lies in a *fidelity-aware comparison protocol* that subjects the same inverse-dynamics PID architecture to multiple independent numerical and physical audits. The specific contributions are:

(i) a two-section reduced Cosserat formulation with configuration-dependent mass, gravity, and Coriolis terms, elastic and viscous generalized forces, and explicit mapping from six bounded tendon tensions to four generalized bending coordinates;

(ii) an inverse-dynamics feedforward plus section-wise PID controller whose six gains are tuned under exactly the same objective and bounds by eleven established metaheuristics, avoiding the confounding effect of a specially privileged "proposed" optimizer;

(iii) dual final-time integration using fourth-order Runge–Kutta (RK4) and implicit Newmark-β, allowing controller rankings to be checked for numerical-scheme sensitivity;

(iv) four robustness cases that separate nominal response, generalized payload loading, measurement noise, and combined payload–impulse–noise exposure;

(v) an explicit full configuration-dependent Cosserat replay of the best reference-linearized controller, revealing a large fidelity gap that would be invisible if evaluation stopped at the tuning model; and

(vi) a reproducible MATLAB package producing ten publication-oriented figures, tuning/performance tables, and structured result files for independent checking.

The rest of the paper is organized as follows. Section 2.1 defines the robot and the fidelity-aware benchmark. Section 2.2 develops the reduced Cosserat kinematics and dynamics. Section 2.3 formulates inverse tendon actuation, inverse-dynamics PID feedback, and the two time integrators. Section 2.4 presents the six-parameter optimization problem and the eleven tuning algorithms. Section 2.5 details the numerical protocol. Section 3 reports the ten figure-based result groups, followed by a broader discussion in Section 4, limitations in Section 4.6, and conclusions in Section 5.

2 Methods

This section presents the aim, design, and computational setting of the study, followed by a complete description of the materials (the reduced Cosserat-rod model and its parameters), the control and numerical-integration methodology, the metaheuristic tuning procedure, and the reproducibility settings used to generate the results reported in Section 3. The study is entirely computational: no human participants, animals, or tissue samples were involved, and no ethics approval or consent statement is therefore required.

2.1 Problem statement and fidelity-aware framework

This section introduces the computational problem before deriving individual equations. The purpose is to make explicit what is optimized, what is held fixed, and which model is used at each stage. This separation is essential because the paper later shows that apparently minor modeling choices can change the numerical conclusion.

The robot consists of two serial continuum sections of equal length $L_1 = L_2 = 0.15$ m. Each section is actuated by three tendons equally spaced at angular locations

$$\gamma_j \in \left\{0, \frac{2\pi}{3}, \frac{4\pi}{3}\right\}, \qquad j = 1,2,3, \tag{1}$$

with tendon-routing radius $r_t = 0.01$ m. Eight spacer disks are represented per section. The resulting robot has six tendon tensions but, under the shear-inextensible and torsion-free reduced strain assumption used here, four generalized bending coordinates,

$$\boldsymbol{q} = \begin{bmatrix} \kappa_{x,1} & \kappa_{y,1} & \kappa_{x,2} & \kappa_{y,2} \end{bmatrix}^{\mathsf{T}} \in \mathbb{R}^4. \tag{2}$$

The controller receives a smooth time-varying reference $\boldsymbol{q}_r(t)$ and must generate tendon tensions that minimize shape and tip-tracking error while remaining inside unilateral tension limits.

The computational framework uses three model levels. During metaheuristic tuning, reference-dependent dynamic quantities are precomputed along the desired trajectory and used as a lookup model. This keeps all eleven optimizers computationally comparable. The

same lookup representation is then used for the main RK4/Newmark benchmark and the four robustness scenarios. Finally, the controller that gives the lowest nominal RK4 tip RMSE is replayed using fully configuration-dependent mass, gravity, and Coriolis terms evaluated at the instantaneous simulated state. This last stage is a *fidelity audit*, not another tuning stage.

The design intentionally separates four questions:

1. **Search question:** which established metaheuristic finds the most effective six PID gains under the common tuning objective?

2. **Integrator question:** does the ranking persist when the same gains are propagated by RK4 and Newmark-β?

3. **Robustness question:** do the stronger controllers retain their advantage under payload-like loads, measurement noise, and a combined disturbance?

4. **Fidelity question:** does the strongest reduced-model controller still perform well when the nonlinear dynamics are recomputed from the actual evolving configuration?

This hierarchy prevents an attractive optimization curve from being treated as sufficient evidence of physical control quality.

2.2 Two-section Cosserat-rod model

The mechanics used in the benchmark are developed in this section. The model is reduced enough to make repeated population-based tuning feasible, but it preserves the central ingredients needed for the later fidelity test: three-dimensional section kinematics, configuration-dependent inertia, gravity, elastic restoring forces, damping, and tendon-generated generalized effort.

2.2.1 Constant-strain section kinematics

For section $i \in \{1,2\}$, the reduced Cosserat strain is written as

$$\boldsymbol{u}_i = \begin{bmatrix} \kappa_{x,i} & \kappa_{y,i} & 0 \end{bmatrix}^{\mathsf{T}}, \qquad \boldsymbol{v} = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^{\mathsf{T}}, \tag{3}$$

where $\boldsymbol{u}_i$ is the bending/twist strain and $\boldsymbol{v}$ imposes unit axial stretch with no shear. With the skew operator $[\boldsymbol{u}]_\times$, the local homogeneous transform over an arc coordinate s is

$$\boldsymbol{g}_i(s) = \begin{bmatrix} \boldsymbol{R}_i(s) & \boldsymbol{p}_i(s) \\ \boldsymbol{0}^\top & 1 \end{bmatrix} = \exp\left(\begin{bmatrix} [\boldsymbol{u}_i]_\times & \boldsymbol{v} \\ \boldsymbol{0}^\top & 0 \end{bmatrix} s \right). \tag{4}$$

Let $\omega_i = \|\boldsymbol{u}_i\|$ and $\theta_i = \omega_i s$. For $\omega_i > 0$, the implementation evaluates

$$\boldsymbol{R}_i(s) = \boldsymbol{I} + \frac{\sin\theta_i}{\omega_i}[\boldsymbol{u}_i]_\times + \frac{1-\cos\theta_i}{\omega_i^2}[\boldsymbol{u}_i]_\times^2, \tag{5}$$

$$\boldsymbol{p}_i(s) = \left[s\boldsymbol{I} + \frac{1-\cos\theta_i}{\omega_i^2}[\boldsymbol{u}_i]_\times + \frac{\theta_i - \sin\theta_i}{\omega_i^3}[\boldsymbol{u}_i]_\times^2 \right] \boldsymbol{v}. \tag{6}$$

The straight configuration follows continuously as $\omega_i \to 0$. Serial composition gives the tip frame

$$\boldsymbol{g}_{\text{tip}}(\boldsymbol{q}) = \boldsymbol{g}_1(L_1)\boldsymbol{g}_2(L_2), \tag{7}$$

and its translation defines the Cartesian tip position $\boldsymbol{p}_{\text{tip}}(\boldsymbol{q}) \in \mathbb{R}^3$.

2.2.2 Distributed mass approximation and generalized inertia

The backbone mass is represented through lumped quadrature points along both sections, augmented by the spacer-disk masses and distal payload. If $\boldsymbol{p}_\ell(\boldsymbol{q})$ is the Cartesian location of lumped mass m_ℓ, its position Jacobian is

$$\boldsymbol{J}_\ell(\boldsymbol{q}) = \frac{\partial \boldsymbol{p}_\ell}{\partial \boldsymbol{q}} \in \mathbb{R}^{3\times 4}. \tag{8}$$

The code evaluates these Jacobians by centered finite differences, and the generalized mass matrix is assembled as

$$\boldsymbol{M}(\boldsymbol{q}) = \sum_{\ell=1}^{N_m} m_\ell \boldsymbol{J}_\ell^\top \boldsymbol{J}_\ell + \varepsilon_M \boldsymbol{I}_4, \tag{9}$$

where ε_M is a small regularization term. This construction retains the configuration dependence induced by the moving distributed masses while avoiding a much larger finite-element state.

The gravitational potential is

$$V_g(\boldsymbol{q}) = \sum_{\ell=1}^{N_m} m_\ell g\, z_\ell(\boldsymbol{q}), \tag{10}$$

and the generalized gravity vector follows from

$$\boldsymbol{G}(\boldsymbol{q}) = \frac{\partial V_g}{\partial \boldsymbol{q}}. \tag{11}$$

Centered finite differences are again used in the implementation. The Coriolis/centrifugal vector is constructed through numerical Christoffel symbols,

$$\Gamma_{ijk}(\boldsymbol{q}) = \frac{1}{2}\left(\frac{\partial M_{ij}}{\partial q_k} + \frac{\partial M_{ik}}{\partial q_j} - \frac{\partial M_{jk}}{\partial q_i}\right), \tag{12}$$

$$c_i(\boldsymbol{q}, \dot{\boldsymbol{q}}) = \sum_{j=1}^{4}\sum_{k=1}^{4} \Gamma_{ijk}\dot{q}_j\dot{q}_k. \tag{13}$$

2.2.3 Elastic and dissipative generalized forces

For the circular backbone with radius $r_b = 2$ mm, the area moment of inertia is

$$I = \frac{\pi r_b^4}{4}, \tag{14}$$

and the two sections use an effective Young's modulus $E_i = 25$ MPa. The bending-energy approximation produces the diagonal generalized stiffness matrix

$$\boldsymbol{K} = \mathrm{diag}\left(EI_1L_1, EI_1L_1, EI_2L_2, EI_2L_2\right). \tag{15}$$

A neutral-configuration mass matrix $\boldsymbol{M}_0 = \boldsymbol{M}(\boldsymbol{0})$ is used to define modal viscous damping. For damping ratio $\zeta_j = 0.10$,

$$D_{jj} = 2\zeta_j\sqrt{M_{0,jj}K_{jj}}, \qquad \boldsymbol{D} = \mathrm{diag}(D_{11}, \ldots, D_{44}). \tag{16}$$

Table 1: Principal robot and dynamic parameters used in the numerical study.

Quantity	Symbol	Value
Number of sections	n_s	2
Section lengths	L_1, L_2	0.15, 0.15 m
Backbone radius	r_b	0.002 m
Tendon/disc radius	r_t, r_d	0.010, 0.010 m
Tendons per section	n_t	3
Disks per section	n_d	8
Effective Young's modulus	E_1, E_2	25 MPa
Backbone density	ρ	1150 kg m^{-3}
Single disk mass	m_d	0.003 kg
Distal payload mass	m_p	0.012 kg
Damping ratio	ζ	0.10
Tendon preload	T_0	2.5 N
Tension interval	$[T_{\min}, T_{\max}]$	[0.25, 16] N

2.2.4 Reduced equations of motion

The resulting four-coordinate dynamic model is

$$\boxed{\boldsymbol{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{c}(\boldsymbol{q}, \dot{\boldsymbol{q}}) + \boldsymbol{D}\dot{\boldsymbol{q}} + \boldsymbol{K}\boldsymbol{q} + \boldsymbol{G}(\boldsymbol{q}) = \boldsymbol{B}\boldsymbol{T} + \boldsymbol{w}_{\text{ext}}.} \tag{17}$$

Here $\boldsymbol{T} \in \mathbb{R}^6$ contains tendon tensions, $\boldsymbol{B} \in \mathbb{R}^{4\times 6}$ is the tendon-to-curvature generalized-force map, and $\boldsymbol{w}_{\text{ext}}$ represents generalized disturbances. Equation (17) is the common mechanical basis for both inverse-dynamics feedforward and forward simulation.

2.3 Inverse tendon actuation, PID control, and time integration

This section connects the Cosserat dynamics to the six physical tendon tensions and then defines the two numerical solvers used for cross-verification. The same controller is deliberately propagated by two different integration philosophies so that numerical sensitivity becomes observable rather than implicit.

2.3.1 Tendon-to-generalized-force mapping

For section i, virtual work gives the 2×3 map

$$\boldsymbol{B}_i = L_i r_t \begin{bmatrix} -\sin\gamma_1 & -\sin\gamma_2 & -\sin\gamma_3 \\ \cos\gamma_1 & \cos\gamma_2 & \cos\gamma_3 \end{bmatrix}. \tag{18}$$

The complete matrix $\boldsymbol{B}$ is block diagonal in the two section groups. Because three tendons generate two bending components per section, the actuation is redundant. Given a desired generalized effort $\boldsymbol{u}_c \in \mathbb{R}^4$, the initial section-level tension estimate is

$$\boldsymbol{T}_i^{(0)} = T_0 \mathbf{1}_3 + \boldsymbol{B}_i^{\dagger} \boldsymbol{u}_{c,i}, \tag{19}$$

followed by projection onto

$$T_{\min} \le T_{i,j} \le T_{\max}. \tag{20}$$

The implementation then performs a small number of residual corrections on the non-saturated tendons. The actually realizable generalized effort is

$$\boldsymbol{u}_a = \boldsymbol{B}\boldsymbol{T}, \qquad \boldsymbol{r}_u = \boldsymbol{u}_c - \boldsymbol{u}_a, \tag{21}$$

where $\boldsymbol{r}_u$ is retained as a diagnostic. This explicit distinction between commanded and realizable effort is important because an unconstrained controller can otherwise appear accurate while requiring physically impossible negative or excessive tendon tensions.

2.3.2 Reference inverse dynamics and section-wise PID

For a desired trajectory $\boldsymbol{q}_r(t)$, its velocity $\dot{\boldsymbol{q}}_r$ and acceleration $\ddot{\boldsymbol{q}}_r$, the reference inverse-dynamics term is

$$\boldsymbol{u}_{\mathrm{ff}} = \boldsymbol{M}_r \ddot{\boldsymbol{q}}_r + \boldsymbol{c}_r + \boldsymbol{D}\dot{\boldsymbol{q}}_r + \boldsymbol{K}\boldsymbol{q}_r + \boldsymbol{G}_r, \tag{22}$$

where the subscript r denotes quantities evaluated along the reference trajectory during the lookup-model stages. Define

$$\boldsymbol{e} = \boldsymbol{q}_r - \boldsymbol{q}, \qquad \dot{\boldsymbol{e}} = \dot{\boldsymbol{q}}_r - \dot{\boldsymbol{q}}. \tag{23}$$

The control command is

$$\boxed{\boldsymbol{u}_c = \boldsymbol{u}_{\mathrm{ff}} + \boldsymbol{K}_p \boldsymbol{e} + \boldsymbol{K}_i \int_0^t \boldsymbol{e}(\tau)\mathrm{d}\tau + \boldsymbol{K}_d \dot{\boldsymbol{e}}.} \tag{24}$$

The two bending coordinates of each section share the same PID gains, giving the six-dimensional search vector

$$\boldsymbol{x} = \begin{bmatrix} K_{p,1} & K_{i,1} & K_{d,1} & K_{p,2} & K_{i,2} & K_{d,2} \end{bmatrix}. \tag{25}$$

This sharing reduces the optimizer dimension while preserving independent tuning of proximal and distal section dynamics.

2.3.3 Fourth-order Runge–Kutta propagation

Writing $\boldsymbol{z} = [\boldsymbol{q}^\top, \dot{\boldsymbol{q}}^\top]^\top$ and $\dot{\boldsymbol{z}} = \boldsymbol{f}(\boldsymbol{z}, t)$, the explicit RK4 update is

$$\boldsymbol{k}_1 = \boldsymbol{f}(\boldsymbol{z}_n, t_n), \tag{26}$$

$$\boldsymbol{k}_2 = \boldsymbol{f}(\boldsymbol{z}_n + \tfrac{\Delta t}{2}\boldsymbol{k}_1, t_n + \tfrac{\Delta t}{2}), \tag{27}$$

$$\boldsymbol{k}_3 = \boldsymbol{f}(\boldsymbol{z}_n + \tfrac{\Delta t}{2}\boldsymbol{k}_2, t_n + \tfrac{\Delta t}{2}), \tag{28}$$

$$\boldsymbol{k}_4 = \boldsymbol{f}(\boldsymbol{z}_n + \Delta t \boldsymbol{k}_3, t_n + \Delta t), \tag{29}$$

$$\boldsymbol{z}_{n+1} = \boldsymbol{z}_n + \frac{\Delta t}{6}(\boldsymbol{k}_1 + 2\boldsymbol{k}_2 + 2\boldsymbol{k}_3 + \boldsymbol{k}_4). \tag{30}$$

RK4 is used during optimization because it provides a deterministic high-order explicit baseline and avoids an inner nonlinear iteration at every objective evaluation.

2.3.4 Implicit Newmark-β propagation

The Newmark method is used as an independent dynamic propagation check. With $\beta = 0.25$ and $\gamma = 0.50$, the predictor is

$$\boldsymbol{q}_{n+1}^{\text{pred}} = \boldsymbol{q}_n + \Delta t \dot{\boldsymbol{q}}_n + \Delta t^2 \left(\frac{1}{2} - \beta \right) \ddot{\boldsymbol{q}}_n, \tag{31}$$

$$\dot{\boldsymbol{q}}_{n+1}^{\text{pred}} = \dot{\boldsymbol{q}}_n + \Delta t (1 - \gamma) \ddot{\boldsymbol{q}}_n. \tag{32}$$

For an acceleration estimate $\ddot{\boldsymbol{q}}_{n+1}$,

$$\boldsymbol{q}_{n+1} = \boldsymbol{q}_{n+1}^{\text{pred}} + \beta \Delta t^2 \ddot{\boldsymbol{q}}_{n+1}, \tag{33}$$

$$\dot{\boldsymbol{q}}_{n+1} = \dot{\boldsymbol{q}}_{n+1}^{\text{pred}} + \gamma \Delta t \ddot{\boldsymbol{q}}_{n+1}. \tag{34}$$

Five fixed-point corrections with relaxation factor 0.72 are used per step. Because Newmark and RK4 process the same control law but resolve the dynamics differently, agreement increases confidence in the result, whereas disagreement is treated as a numerical-sensitivity finding.

2.4 Metaheuristic PID tuning problem

The optimization layer is described here before presenting results. All eleven methods receive the same gain bounds, population size, iteration budget, initial random-seed policy, dynamic model, and scalar objective. Consequently, differences in the tuning outcome are attributable to the search dynamics rather than to unequal problem formulations.

2.4.1 Multi-term objective function

The tuning cost balances normalized curvature error, time-weighted error, peak error, generalized actuation effort, and tendon saturation. Let $q_s = 5\ \text{m}^{-1}$ be the curvature normalization and

$u_s = 8 \times 10^{-4}$ the generalized-effort normalization. The components are

$$J_{\mathrm{MSE}} = \frac{1}{4N} \sum_{k=1}^{N} \left\| \frac{\boldsymbol{e}_k}{q_s} \right\|_2^2, \tag{35}$$

$$J_{\mathrm{ITAE}} = \frac{1}{t_f^2 q_s} \int_0^{t_f} t \, \|\boldsymbol{e}(t)\|_2 \mathrm{d}t, \tag{36}$$

$$J_{\max} = \frac{1}{q_s} \max_t \|\boldsymbol{e}(t)\|_2, \tag{37}$$

$$J_u = \frac{1}{4N} \sum_{k=1}^{N} \left\| \frac{\boldsymbol{u}_{a,k}}{u_s} \right\|_2^2, \tag{38}$$

$$J_{\mathrm{sat}} = \frac{1}{N} \sum_{k=1}^{N} s_k, \tag{39}$$

where s_k is the fraction of the six tendons at a tension bound. The scalar objective used by every optimizer is exactly

$$\boxed{J(\boldsymbol{x}) = 0.46 J_{\mathrm{MSE}} + 0.24 J_{\mathrm{ITAE}} + 0.10 J_{\max} + 0.12 J_u + 0.08 J_{\mathrm{sat}}.} \tag{40}$$

A diverged or nonfinite simulation is assigned a large penalty. The weighting deliberately gives tracking error the largest influence while preventing a search algorithm from obtaining an attractive error by systematically driving tendons against their limits.

2.4.2 Eleven optimizers under a common budget

The comparison includes PSO, GA, ABC, continuous ACO (ACOR style), BOA, WOA, GWO, TLBO, HHO, POA, and CS. No hybrid or specially enhanced method is inserted into the benchmark. This choice is important for two reasons. First, the paper focuses on the control/modeling framework rather than claiming novelty from a renamed optimizer. Second, using recognizable baselines makes the later ranking inversions easier to diagnose because each search process is independently established.

For compactness, each method can be viewed as generating candidate gains by a characteristic exploration–exploitation operator. PSO uses inertial velocity with personal/global attraction; GA uses tournament selection, arithmetic crossover, and mutation; ABC alternates employed, onlooker, and scout phases; ACOR samples Gaussian kernels around an ordered archive; BOA

Algorithm 1 Fidelity-aware inverse-dynamics PID tuning and evaluation pipeline

Require: Robot parameters $\mathscr{P}$, reference trajectory $\boldsymbol{q}_r(t)$, optimizer set $\mathscr{A}$
Ensure: Tuned controllers, dual-integrator results, robustness results, full-model fidelity audit

```
Assemble K, neutral mass M_0, damping D, and tendon map B
Precompute reference lookup dynamics {M_r, c_r, G_r, u_ff}
for each optimizer a ∈ 𝒜 do
    Initialize a population of six-gain vectors x
    for k = 1 to N_iter do
        Simulate the tuning trajectory with RK4 and bounded tendon inverse actuation
        Evaluate J(x) from Eq. (40)
        Update the population using optimizer a
    end for
    Retain the lowest-cost gain vector for optimizer a
end for
for each tuned controller do
    Run final nominal simulation with RK4
    Run final nominal simulation with Newmark-β
    Run nominal, payload, noise, and combined robustness cases
end for
Select the lowest nominal RK4 tip-RMSE controller
Replay that controller with configuration-dependent M(q), c(q, q̇), and G(q)
Compute tracking, effort, saturation, residual, and fidelity-gap metrics
```

uses fragrance-weighted global/local moves; WOA alternates encircling and spiral motion; GWO combines alpha, beta, and delta leader estimates; TLBO alternates teacher and learner phases; HHO uses escaping-energy-dependent exploration/exploitation; POA uses prey-directed and local-radius phases; and CS combines Levy flights with nest abandonment. Candidate gains are clipped to the same physically scaled bounds after each update.

2.5 Numerical protocol and reproducibility settings

This section specifies the simulation settings needed to reproduce the reported numbers. It also distinguishes exploratory stochastic settings from statistically replicated optimization, because the supplied run uses one optimization run per method. That design is sufficient for a deterministic benchmark snapshot but does not support claims of statistical superiority between stochastic optimizers.

The tuning trajectory is simulated for 4.8 s with $\Delta t = 0.04$ s, while the final comparison uses 8.0 s with $\Delta t = 0.02$ s. Each optimizer uses population size 18 and 24 iterations. A fixed seed policy makes the supplied one-run comparison exactly reproducible. The reference curvature

Algorithm 2 Bounded tendon inverse actuation for one control sample

Require: Generalized command $\boldsymbol{u}_c$, preload T_0, bounds $T_{\min}, T_{\max}$

```
for section i = 1,2 do
    T_i ← T_0 1 + B_i^† u_{c,i}
    Project T_i to [T_min, T_max]
    for r = 1 to 4 do
        e_u ← u_{c,i} − B_i T_i
        Identify tendons that are not saturated
        if at least one tendon is free then
            Correct free tensions with the pseudoinverse of the free-column submatrix
            Re-project to [T_min, T_max]
        else
            break
        end if
    end for
end for
Return T, u_a = BT, residual r_u = u_c − u_a
```

trajectory is

$$\kappa_{x,1}(t) = 4.4\sin(0.60t), \tag{41}$$

$$\kappa_{y,1}(t) = 3.2\sin(0.42t + 0.55), \tag{42}$$

$$\kappa_{x,2}(t) = 3.9\sin(0.60t + 0.95) + 0.75\sin(0.92t), \tag{43}$$

$$\kappa_{y,2}(t) = 3.4\sin(0.47t - 0.35), \tag{44}$$

with analytic first and second derivatives used in the inverse-dynamics term.

The robustness protocol comprises four cases. In the payload case, a generalized load is applied from $t = 3.1$ s onward. The noise case perturbs curvature and curvature-rate measurements with zero-mean Gaussian noise of standard deviations 0.055 m^{-1} and 0.12 $\mathrm{m}^{-1}\mathrm{s}^{-1}$, respectively. The combined case adds the payload, measurement noise, and a generalized impulse acting from 5.25 to 5.55 s. All methods are tested on identical generated disturbances within a scenario.

Table 2: Numerical and optimization settings.

Setting	Symbol	Value
Tuning time step	Δt_{tune}	0.04 s
Tuning horizon	$t_{f,\text{tune}}$	4.8 s
Final time step	Δt_{final}	0.02 s
Final horizon	$t_{f,\text{final}}$	8.0 s
Population size	N_p	18
Iterations	N_{iter}	24
Runs per optimizer in supplied benchmark	N_r	1
Newmark parameters	(β, γ)	(0.25, 0.50)
Newmark fixed-point iterations	–	5
Newmark relaxation	–	0.72
Payload start	t_p	3.1 s
Impulse interval	–	5.25–5.55 s
Curvature noise std.	σ_q	0.055 m^{-1}
Curvature-rate noise std.	$\sigma_{\dot{q}}$	0.12 $\text{m}^{-1}\text{s}^{-1}$

The main reported metrics are

$$\text{RMSE}_p = 1000\sqrt{\frac{1}{N}\sum_{k=1}^{N}\|\boldsymbol{p}_k - \boldsymbol{p}_{r,k}\|_2^2} \quad [\text{mm}], \tag{45}$$

$$\text{RMSE}_\kappa = \sqrt{\frac{1}{4N}\sum_{k=1}^{N}\|\boldsymbol{q}_k - \boldsymbol{q}_{r,k}\|_2^2}, \tag{46}$$

$$\mathscr{E}_{T^2} = \int_0^{t_f}\sum_{j=1}^{6} T_j^2(t)\mathrm{d}t, \tag{47}$$

$$R_u = \sqrt{\frac{1}{4N}\sum_{k=1}^{N}\|\boldsymbol{r}_{u,k}\|_2^2}, \tag{48}$$

$$S = 100\,\frac{1}{N}\sum_{k=1}^{N} s_k. \tag{49}$$

The quantity $\mathscr{E}_{T^2}$ is explicitly referred to as a *quadratic tendon-effort proxy*; it is not claimed to be electrical or thermodynamic energy.

3 Results

This section reports the numerical outcomes in the same logical order as the ten generated figures. Each subsection first states the question addressed by the corresponding figure and then

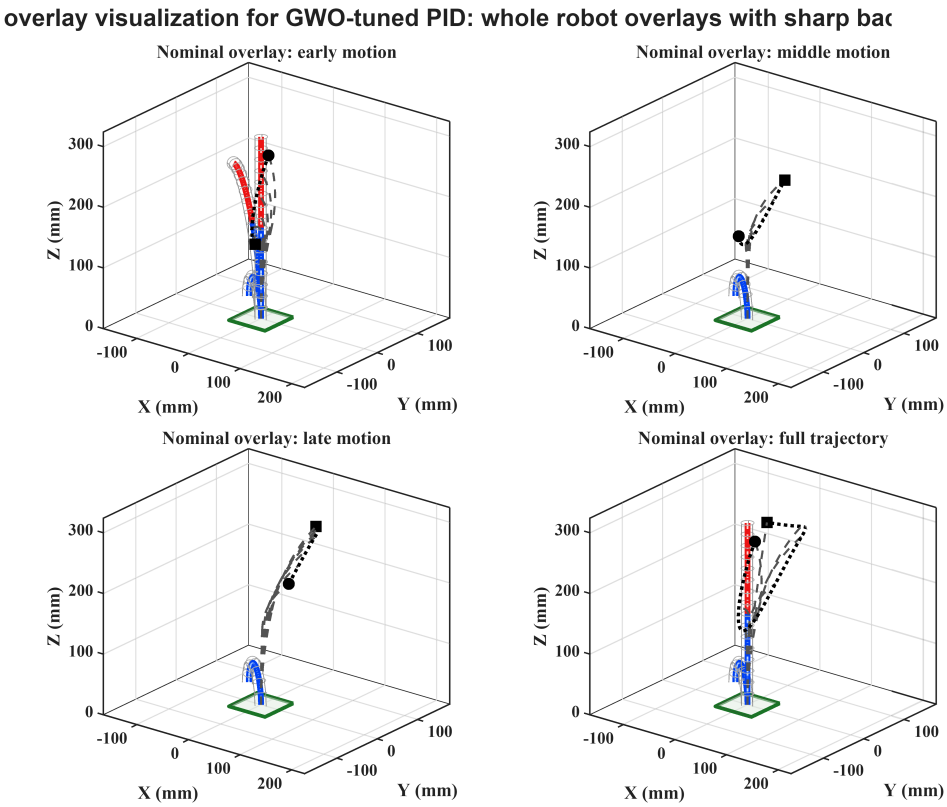

Figure 1: Full configuration-dependent Cosserat replay of the GWO-tuned inverse-dynamics PID controller. Each panel overlays multiple complete robot poses and sharp backbone traces over a different portion of the motion. The spatial separation between actual and reference paths becomes substantial as the motion evolves, exposing the model-fidelity gap.

connects the visual trends to quantitative metrics. Because the optimizer budget contains one run per method in the supplied dataset, terms such as “best” and “stronger” refer to the observed deterministic realization and are not used as statistical claims.

3.1 Whole-robot motion and the full-model fidelity audit

The first result asks whether the selected controller preserves its apparent tracking quality when it is removed from the reference-linearized lookup plant and replayed with configuration-dependent Cosserat dynamics. GWO is selected because it gives the lowest nominal RK4 tip RMSE in the lookup benchmark. Figure 1 overlays multiple complete robot poses, including disks and tendons, together with sharp backbone traces and the reference motion over early, middle, late, and full-trajectory intervals.

The overlays show that the robot remains dynamically bounded and physically coherent, but the reference path and actual backbone diverge markedly during the middle and late phases.

Table 3: Fidelity audit for the GWO-tuned controller. The full model recomputes configuration-dependent inertial, Coriolis, and gravity terms at the simulated state.

Metric	Reference-linearized RK4	Full Cosserat replay
Tip RMSE (mm)	24.699	400.589
Maximum tip error (mm)	–	436.904
Curvature RMSE (m^{-1})	0.490	11.809
Quadratic tendon-effort proxy	391.063	32.052
Tension saturation (%)	2.660	0.457
Peak tension (N)	5.801	6.440
Inverse-actuation residual RMS	3.33×10^{-19}	9.25×10^{-20}

Quantitatively, the same GWO gains that achieve 24.699 mm nominal tip RMSE on the reference-linearized RK4 model produce 400.589 mm tip RMSE in the full configuration-dependent replay. The maximum tip error reaches 436.904 mm, while curvature RMSE rises from 0.490 to 11.809 m^{-1}. This is a factor of approximately 16.2 in tip RMSE. The result demonstrates that a controller can be excellent relative to a reduced benchmark yet fail to transfer when inertial, gravity, and Coriolis quantities are recomputed at the actual configuration.

This finding gives the benchmark its main methodological significance. If evaluation ended after optimizer convergence or even after the nominal lookup simulation, GWO would appear decisively successful. The full replay instead reveals that the lookup approximation is optimistic outside the reference neighborhood. The result does not invalidate reduced-order tuning; it identifies the condition under which reduced-order tuning needs either iterative re-linearization, robustification, gain scheduling, online adaptation, or a full-model-in-the-loop correction stage.

3.2 Geometric response under nominal and disturbed scenarios

The second figure examines whether the selected controller's overall deformation pattern changes visibly under the four benchmark scenarios. In contrast to Figure 1, these scenario runs use the common lookup-model representation so that the disturbance effect can be compared without changing model fidelity at the same time.

The four panels retain a similar gross bending envelope, indicating that the GWO controller does not lose boundedness under the imposed perturbations. The most visible changes occur in the later portions of payload and combined trajectories, where the backbone family spreads

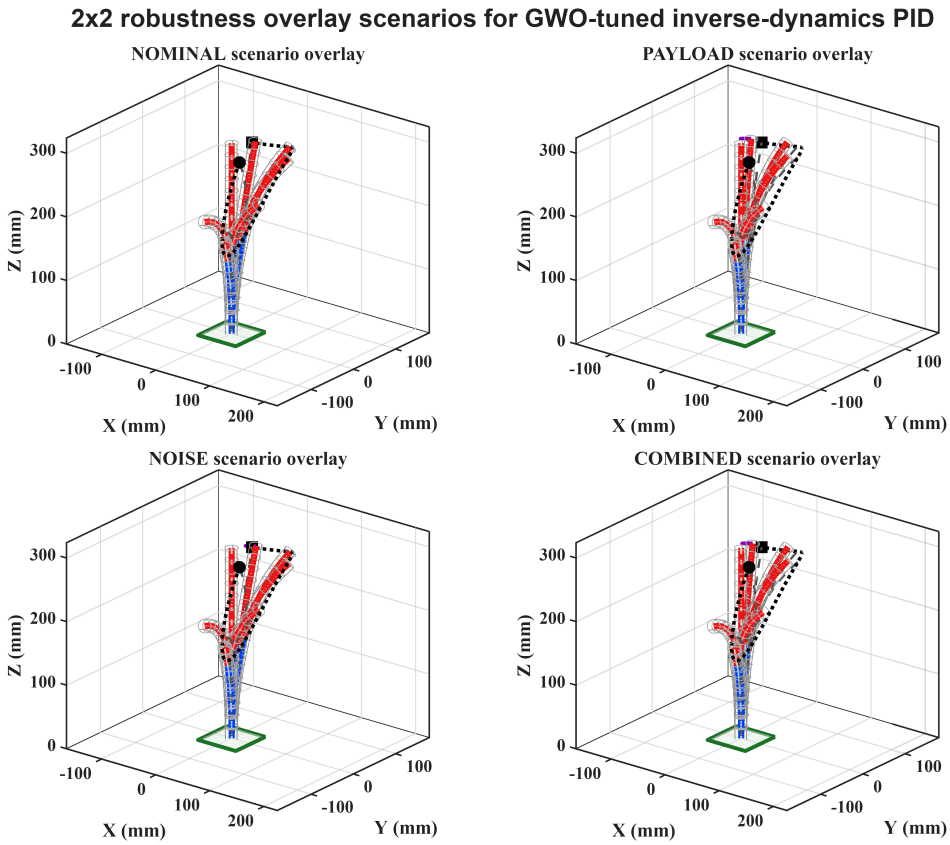

Figure 2: Two-by-two whole-robot overlays for the GWO-tuned controller under nominal, payload, measurement-noise, and combined disturbance conditions using the reference-linearized plant representation. Multiple disk, tendon, and backbone configurations are superimposed in each panel.

relative to the nominal envelope. The quantitative results support this reading: the GWO tip RMSE changes from 24.699 mm nominally to 27.523 mm with payload, 25.055 mm with measurement noise, and 27.265 mm in the combined case. Thus the combined disturbance raises error by roughly 10.4% relative to nominal, which is modest compared with the much larger model-fidelity gap in Table 3. In other words, for this benchmark the modeling approximation is a more severe threat than the specific disturbance amplitudes that were applied.

3.3 Cartesian tip tracking across the eleven tuners

The third result compares the practical consequence of metaheuristic gain selection in Cartesian space. Figure 3 plots x, y, and z tip coordinates together with the norm of the tip error for all eleven tuned PID controllers.

A clear two-group behavior appears. GWO, POA, WOA, TLBO, and HHO follow the

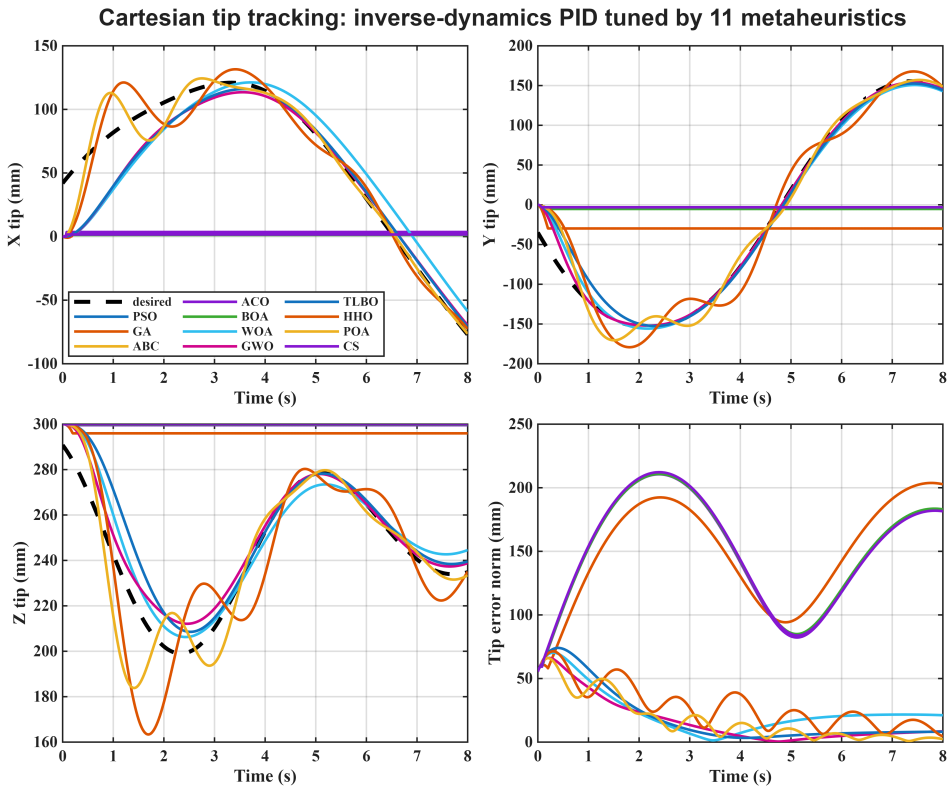

Figure 3: Cartesian tip tracking for the eleven metaheuristic-tuned inverse-dynamics PID controllers under nominal RK4 simulation. The fourth panel reports the instantaneous Euclidean tip-error norm.

reference substantially more closely than PSO, GA, ABC, ACO, BOA, and CS in the RK4 benchmark. The best two RK4 tip RMSE values are 24.699 mm for GWO and 24.945 mm for POA, followed by TLBO at 29.214 mm and WOA at 29.413 mm. HHO reaches 33.304 mm. In contrast, GA, ACO, ABC, PSO, BOA, and CS fall in a narrow poor-performing band of approximately 154.7–156.6 mm.

This result is especially important because the tuning objective values themselves differ only slightly. GWO's best tuning cost is 0.414666, while ACO, which yields 155.617 mm RK4 tip RMSE, obtains a nearly identical cost of 0.414872. A change of only 2.06×10^{-4} in the scalar objective therefore corresponds to more than a six-fold difference in final tip RMSE. The objective combines curvature error, time weighting, effort, and saturation on a shorter tuning trajectory; consequently, its numerical value should be treated as a selection criterion rather than as a direct surrogate for final Cartesian accuracy.

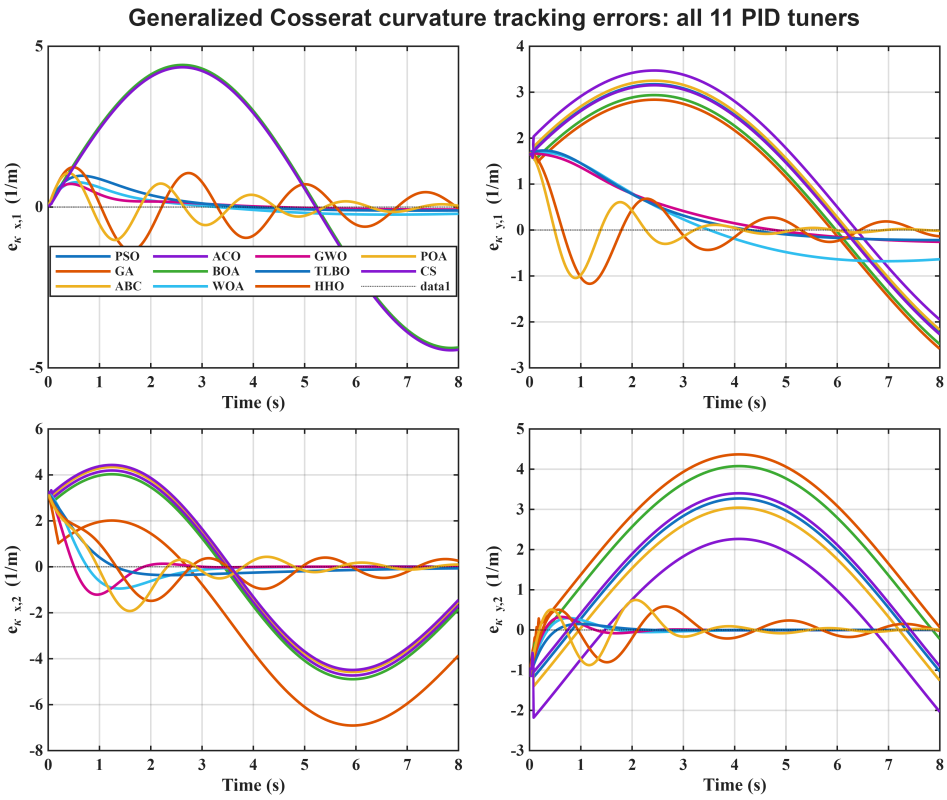

Figure 4: Tracking errors for $\kappa_{x,1}$, $\kappa_{y,1}$, $\kappa_{x,2}$, and $\kappa_{y,2}$ under nominal RK4 simulation.

3.4 Generalized curvature tracking errors

The fourth figure examines the same controller comparison in the native generalized coordinates of the reduced Cosserat model. This view distinguishes whether tip-error differences originate from one section or from systematic tracking failure across all four bending components.

The generalized-coordinate curves confirm the split seen in Cartesian space. GWO produces the lowest curvature RMSE, 0.490 m^{-1}, followed by POA at 0.541, TLBO at 0.553, WOA at 0.572, and HHO at 0.645 m^{-1}. The remaining six methods exceed 2.72 m^{-1}. Because the tip position is obtained through serial exponentials of the section strains, sustained curvature errors accumulate geometrically and explain the much larger Cartesian deviations of the weaker group.

The comparison also illustrates why evaluating only the tip can be incomplete. Two controllers can show similar tip error at a given instant while distributing curvature differently between sections, which can change tendon loads, collision clearance, and structural stress. The four-coordinate error therefore provides a shape-level diagnostic that complements endpoint accuracy.

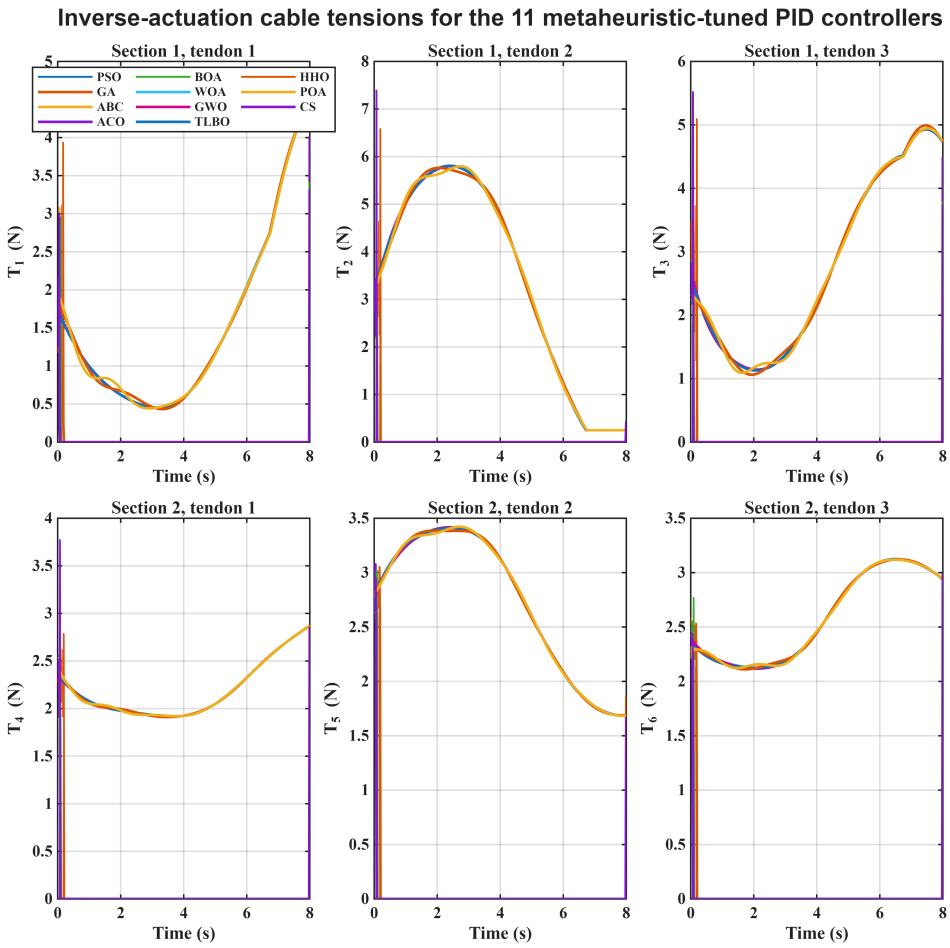

Figure 5: Six tendon-tension histories for the eleven tuned PID controllers. Tensions include the common preload and are projected onto the admissible interval before the generalized actuation is applied.

3.5 Tendon tensions and feasibility

Figure 5 moves from geometric performance to the physical actuation variables. The six panels show the tension histories for each tendon and each controller. Because tendon actuation is unilateral, these trajectories are more informative than an unconstrained generalized torque alone.

The high-tracking group uses substantially more dynamic tendon modulation than the poor-tracking group. For GWO, the nominal quadratic tendon-effort proxy is 391.063 and the peak tension is approximately 5.801 N, well below the 16 N upper limit. GWO's average saturation fraction corresponds to 2.660% of samples/tendons at a bound. POA shows a closely related pattern with effort proxy 390.738 and saturation 2.618%. By comparison, several weak controllers produce effort proxies between approximately 3.5 and 9.3, indicating that their low

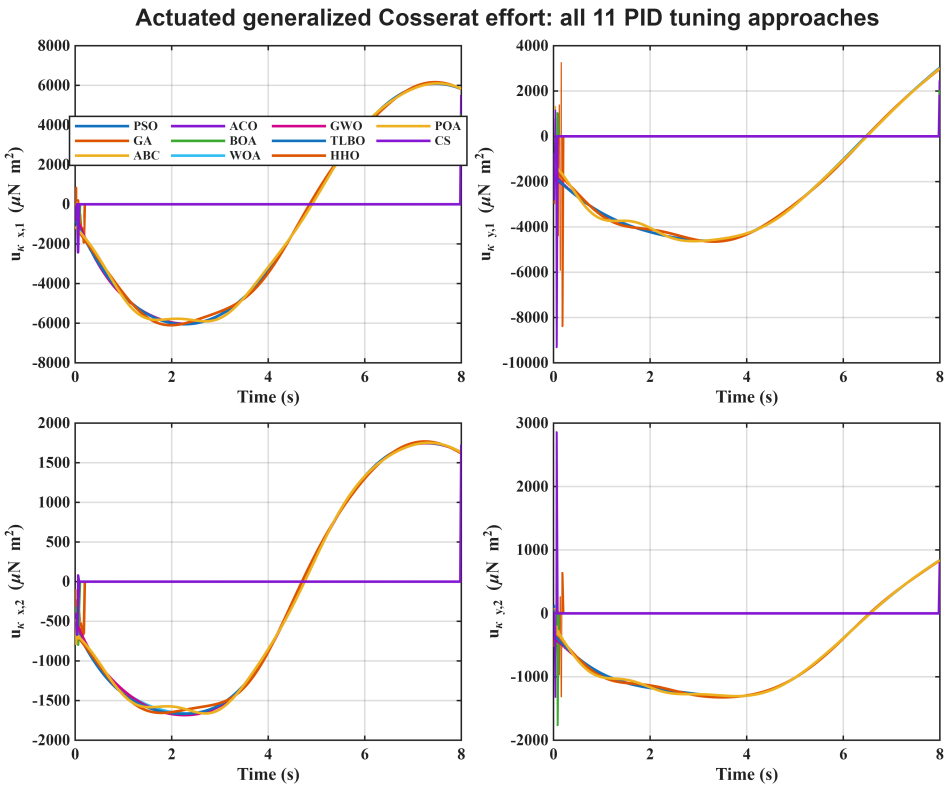

Figure 6: Realized generalized bending effort for the four Cosserat curvature coordinates. The curves are the post-saturation actuation $\boldsymbol{u}_a = \boldsymbol{BT}$, not the unconstrained PID command.

actuation demand is not evidence of superior efficiency; rather, they fail to generate enough corrective deformation to follow the reference.

This trade-off is important when reading optimization results. A very low effort metric can be achieved by a controller that simply does not track. Conversely, a large quadratic tension integral is not automatically undesirable if it remains within feasible tension limits and is required for accurate motion. For this reason, the present objective includes effort but does not let it dominate tracking error.

3.6 Generalized control effort

The sixth figure shows the four realized generalized effort components $\boldsymbol{u}_a = \boldsymbol{BT}$. This representation removes tendon redundancy and reveals the net bending actions supplied to the reduced coordinates.

The strong controllers exhibit larger, phase-dependent generalized efforts that mirror the oscillatory reference curvature. The weak group remains comparatively close to low-effort

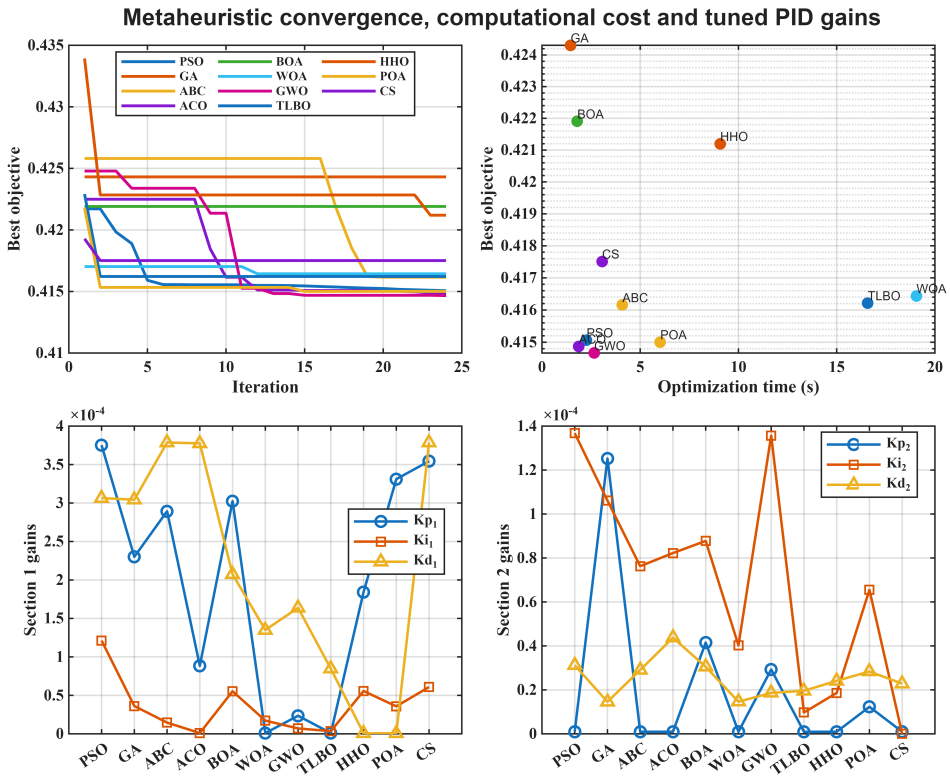

Figure 7: Metaheuristic tuning results: convergence histories, objective-versus-runtime behavior, and the optimized section-wise PID gains. All methods use the same population size and iteration budget.

trajectories and therefore cannot compensate the combined elastic, damping, gravity, and dynamic terms. An important numerical check is that the inverse-actuation residual remains extremely small: nominal residual RMS values are on the order of 10^{-20}–10^{-19} for all methods. Thus the major tracking differences do not arise from a failure of the tendon pseudoinverse mapping. They arise from the gain choices and their dynamic interaction with the plant.

3.7 Optimizer convergence, runtime, and section-wise PID gains

The seventh figure compares search behavior rather than closed-loop time histories. This view is useful because it reveals that similar objective values can be reached with very different computational costs and gain combinations.

GWO attains the smallest observed tuning cost, 0.414666, followed closely by ACO (0.414872), POA (0.415006), and PSO (0.415065). These differences are small relative to the large spread in final RK4 tip error, again showing that optimization cost is not a complete performance descriptor. Runtime differs much more strongly. GA is the fastest in the recorded

Table 4: Observed optimizer tuning outcome. Scientific notation is used because the generalized curvature gains are small in SI units.

Method	J^*	Time (s)	K_{p1}	K_{i1}	K_{d1}	K_{p2}	K_{i2}	K_{d2}
PSO	0.415065	2.267	3.752e-4	1.21e-4	3.064e-4	9.425e-7	1.37e-4	3.10e-5
GA	0.424308	1.457	2.301e-4	3.60e-5	3.043e-4	1.252e-4	1.06e-4	1.50e-5
ABC	0.416157	4.085	2.893e-4	1.50e-5	3.787e-4	9.425e-7	7.60e-5	2.90e-5
ACO	0.414872	1.870	8.835e-5	1.00e-6	3.775e-4	9.687e-7	8.20e-5	4.40e-5
BOA	0.421907	1.794	3.025e-4	5.50e-5	2.074e-4	4.150e-5	8.80e-5	3.10e-5
WOA	0.416438	19.061	9.425e-7	1.70e-5	1.350e-4	9.425e-7	4.00e-5	1.50e-5
GWO	**0.414666**	2.657	2.352e-5	7.00e-6	1.638e-4	2.928e-5	1.36e-4	1.90e-5
TLBO	0.416217	16.574	9.425e-7	3.00e-6	8.460e-5	9.425e-7	1.00e-5	2.00e-5
HHO	0.421196	9.068	1.843e-4	5.60e-5	7.574e-7	9.425e-7	1.90e-5	2.40e-5
POA	0.415006	6.019	3.311e-4	3.60e-5	7.574e-7	1.229e-5	6.60e-5	2.80e-5
CS	0.417511	3.060	3.544e-4	6.10e-5	3.787e-4	9.425e-7	0	2.30e-5

run at 1.457 s, followed by BOA at 1.794 s and ACO at 1.870 s. WOA and TLBO require 19.061 and 16.574 s, respectively, while HHO requires 9.068 s.

The optimized gains also reveal that multiple parameter combinations can yield almost indistinguishable tuning costs. Distal-section proportional gains frequently approach the lower bound, whereas integral and derivative gains redistribute across methods. Such non-uniqueness is expected in a redundant, nonlinear problem with a scalarized objective. It also explains why ranking by J alone can be fragile: different gain vectors can trade transient shape error against actuation effort and then behave differently on the longer final trajectory.

3.8 Robustness across payload, noise, and combined disturbances

Figure 8 evaluates whether the nominal ranking is preserved when disturbances are introduced. This is a more meaningful robustness check than simply perturbing the optimizer because the same tuned gains are subjected to altered plant inputs and measurement quality.

POA and GWO remain the strongest controllers over the disturbance suite. Under payload loading, POA gives 24.989 mm tip RMSE and GWO gives 27.523 mm. Under measurement noise the two methods are almost tied at 25.058 and 25.055 mm. In the combined case, POA gives the lowest observed error at 25.062 mm, while GWO reaches 27.265 mm. WOA remains below 32 mm in the combined condition, whereas TLBO deteriorates more noticeably to 37.436 mm and HHO reaches 35.371 mm. The weaker group remains near 155–157 mm, indicating that the disturbance does not rescue poorly tuned controllers.

The robustness plot therefore supports two conclusions. First, the nominal high-performing

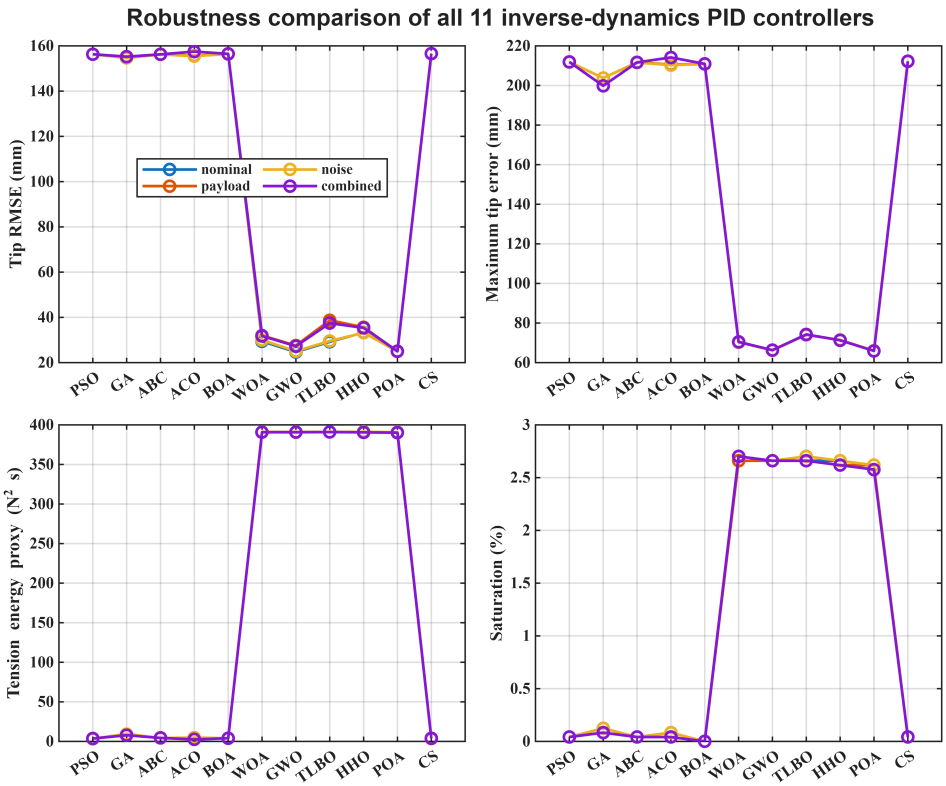

Figure 8: Robustness comparison across nominal, payload, measurement-noise, and combined disturbance scenarios. Metrics include tip tracking, curvature tracking, actuation burden, and saturation behavior.

cluster is not a single-trajectory artifact; its advantage persists under the specified perturbations. Second, the relative ranking within that cluster is scenario dependent: GWO is best nominally and under measurement noise, whereas POA is best under payload and combined disturbance. This suggests that robust controller selection should consider a scenario ensemble instead of relying on a single nominal score.

3.9 RK4 versus Newmark-β: ranking stability and numerical sensitivity

The ninth figure addresses a question that is often omitted in controller-comparison papers: would the same controller appear good if the dynamic equations were advanced by a different standard integration scheme? Figure 9 compares tip and curvature metrics for RK4 and Newmark-β using identical tuned gains.

For GWO, WOA, and TLBO the two integrators agree closely: GWO gives 24.699 mm with RK4 and 24.836 mm with Newmark, WOA gives 29.413 and 29.451 mm, and TLBO gives 29.214 and 29.283 mm. This agreement strengthens confidence that their reduced-model

Table 5: Tip RMSE (mm) for the five strongest nominal RK4 controllers across the four robustness scenarios.

Method	Nominal	Payload	Noise	Combined
GWO	**24.699**	27.523	**25.055**	27.265
POA	24.945	**24.989**	25.058	**25.062**
TLBO	29.214	38.715	29.558	37.436
WOA	29.413	31.994	29.988	31.796
HHO	33.304	35.626	33.198	35.371

Table 6: Nominal performance of all eleven tuned controllers under RK4 and Newmark-β.

Method	RK4 tip RMSE (mm)	Newmark tip RMSE (mm)	Curv. RMSE (m^{-1})	$\mathscr{E}_{T^2}$ –	Sat. (%) –	R_u –
PSO	156.304	156.847	2.773	3.522	0.042	3.99e-20
GA	154.673	16.902	3.236	9.324	0.125	4.11e-20
ABC	156.259	156.520	2.754	4.186	0.042	4.25e-20
ACO	155.617	158.912	2.726	4.616	0.083	5.93e-20
BOA	156.409	**15.990**	2.907	3.978	0.000	2.70e-20
WOA	29.413	29.451	0.572	391.136	2.660	3.19e-19
GWO	**24.699**	24.836	**0.490**	391.063	2.660	3.33e-19
TLBO	29.214	29.283	0.553	391.129	2.660	3.37e-19
HHO	33.304	31.646	0.645	391.231	2.660	3.22e-19
POA	24.945	23.905	0.541	390.738	2.618	3.24e-19
CS	156.570	157.370	2.786	3.786	0.042	3.85e-20

performance is not an RK4-specific artifact. POA differs slightly more, improving from 24.945 mm in RK4 to 23.905 mm in Newmark, but the ranking remains strong.

Two dramatic exceptions occur. GA changes from 154.673 mm under RK4 to only 16.902 mm under Newmark, while BOA changes from 156.409 to 15.990 mm. Those inversions are too large to dismiss as numerical rounding. They indicate that the corresponding closed loops operate in a regime where integration properties, implicit correction, or discrete feedback timing strongly affect the computed trajectory. ACO shows a smaller but still visible change in the opposite direction, from 155.617 to 158.912 mm. The practical consequence is that reporting a single integrator could lead to the strong claim that BOA or GA is among the best methods, whereas RK4 reaches the opposite conclusion.

3.10 Inverse-dynamics and tendon-actuation consistency

The tenth figure consolidates diagnostics that determine whether a tracking result is produced by a numerically consistent actuation solution. This final result group is important because good tracking is not useful if it is obtained with excessive residual, pathological saturation, or hidden violations of the tendon model.

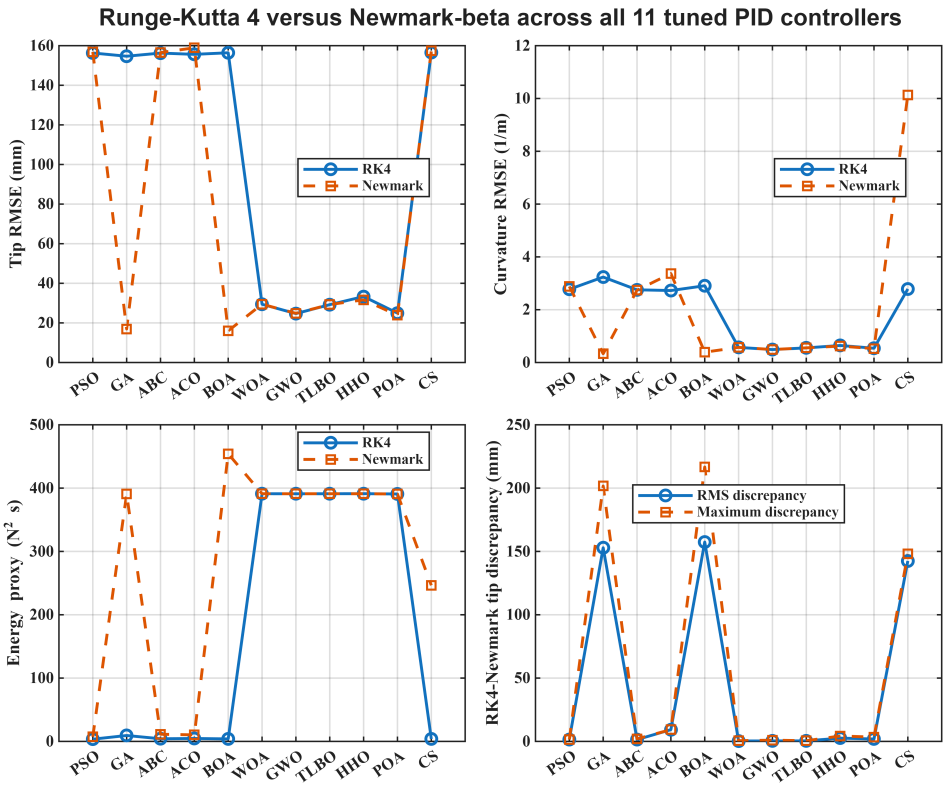

Figure 9: Cross-verification of all eleven tuned controllers using explicit RK4 and implicit Newmark-β integration. Large discrepancies identify controller/integrator combinations that are numerically sensitive.

All inverse-actuation residual RMS values are extremely small, confirming that the pseudoinverse-plus-correction stage reproduces the realizable generalized commands to numerical precision. Peak tensions remain well within the 16 N upper limit. The main feasibility cost in the strong-tracking group is a small saturation percentage of approximately 2.6%, not a wholesale collapse of the tension solution. The figure therefore rules out a simple alternative explanation for the ranking: GWO and POA are not strong merely because the inverse actuation silently ignores constraints; their tendon solutions remain bounded and nearly residual-free.

The combined view also exposes the central control trade-off. Strong lookup-model tracking requires much higher quadratic tendon effort than weak tracking, while full-model replay lowers that effort only because the controller no longer follows the desired motion effectively. Consequently, “lower actuation” must always be read together with tracking quality. A physically useful controller lies on a joint accuracy–feasibility–effort surface rather than at the minimum of any single metric.

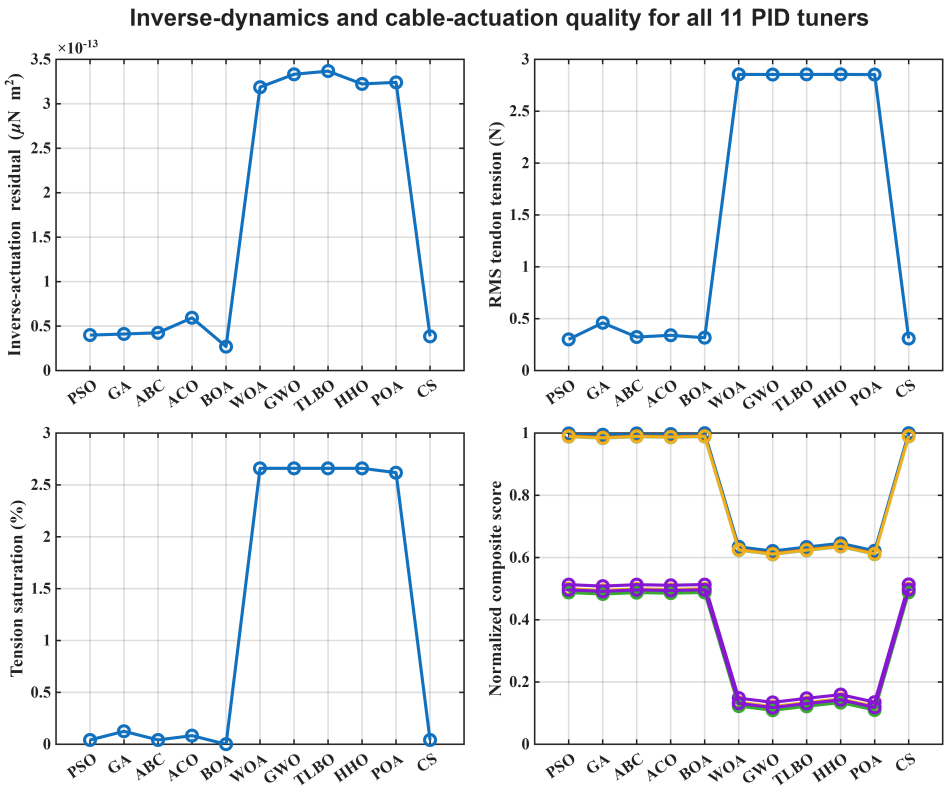

Figure 10: Summary of inverse-dynamics and tendon-actuation quality for the eleven controllers, including residual, peak tension, saturation, and normalized performance/effort indicators.

4 Discussion

The numerical study yields several broader lessons that extend beyond the particular two-section geometry. This section consolidates those lessons and explains why the benchmark is useful even though it does not introduce a new metaheuristic.

4.1 Optimizer ranking is a closed-loop property, not an objective-value property

The first important result is the weak correspondence between the scalar tuning cost and the final Cartesian ranking. GWO, ACO, POA, and PSO all obtain objective values around 0.415, yet their final RK4 tip RMSE values range from 24.699 mm to more than 156 mm. This disparity is caused by three layers of mismatch between J and final Cartesian performance: the tuning horizon is shorter than the final horizon; the objective is defined primarily in generalized curvature coordinates; and effort/saturation penalties trade against error. A metaheuristic therefore should

not be declared superior solely because its terminal cost is lower by a small amount.

The practical recommendation is to separate *optimization metrics* from *validation metrics*. The optimization metric can be compact and inexpensive, but the validation stage should include Cartesian error, shape error, actuator feasibility, robustness, and independent integration. This distinction is particularly relevant in continuum robotics because small curvature biases can integrate into large endpoint displacements.

4.2 Integrator cross-verification can expose hidden instability or discrete sensitivity

The GA and BOA ranking inversions are among the most striking outcomes. A method that appears unusable under RK4 becomes the best or second-best under Newmark. Such a change cannot be explained by metaheuristic randomness because the same gains are reused. It arises after tuning, during dynamic propagation. Several mechanisms may contribute: Newmark's implicit correction can introduce effective numerical damping; the control update is evaluated at different discrete states; saturation and integral-state evolution may interact differently with the predictor–corrector loop; and a gain set near a stability boundary can amplify small integration differences.

This observation argues for a simple but powerful reporting practice: at least one independent integration scheme should be used when nonlinear flexible-robot control conclusions are sensitive to simulation. Agreement does not prove physical validity, but disagreement provides immediate evidence that the conclusion is numerically fragile. In the present benchmark, the GWO/POA/WOA/TLBO group is more credible precisely because most of those controllers preserve similar performance across RK4 and Newmark.

4.3 The largest risk is model fidelity, not the tested disturbance amplitude

The robustness cases change GWO tip RMSE by only a few millimeters relative to nominal, whereas switching from the reference-linearized lookup model to full configuration-dependent dynamics increases the same metric by hundreds of millimeters. This hierarchy is instructive.

It would be easy to emphasize “robustness to noise and payload” while overlooking that the modeling approximation itself is a much stronger disturbance.

The full replay therefore changes the scientific message. The paper does not claim that the current GWO-tuned PID is ready for deployment. Instead, it establishes a reproducible procedure that can reveal when a controller tuned on an efficient reduced model is not transferable. That failure is valuable because it identifies where the next methodological effort should be invested: fidelity-aware tuning, adaptive inverse dynamics, online parameter estimation, robust gain scheduling, or inclusion of full-model samples in the optimizer objective.

4.4 Why the integrated benchmark is useful

The framework combines elements that are often studied separately: Cosserat kinematics, inverse dynamics, bounded tendon allocation, PID tuning, metaheuristic comparison, dual integration, scenario robustness, and full-model replay. None of these ingredients alone is unprecedented. Their combination creates a stricter validation chain in which every attractive result must survive several independent checks.

This integrated perspective is particularly useful for research reproducibility. A reader can reproduce the optimization table, inspect the six tendon tensions, compare RK4 and Newmark, rerun disturbances, and then switch the plant fidelity without changing the controller. That makes the benchmark suitable as a testbed for future control ideas. For example, a neural controller, MPC method, adaptive law, or new optimizer could be inserted into the same evaluation sequence and judged on exactly the same fidelity and numerical criteria.

4.5 Implications for future controller design

The fidelity audit suggests that the most promising next step is not to intensify the metaheuristic search around the same lookup model. A more sophisticated optimizer cannot remove model mismatch that is absent from its objective. More productive strategies include multi-fidelity optimization in which cheap lookup simulations are periodically corrected by full-model evaluations; robust objectives that sample perturbed inertial and stiffness parameters; scheduled inverse-dynamics tables indexed by the actual configuration; and online residual learning that

estimates the discrepancy between the reference model and the full dynamics.

A second direction is to elevate tendon variables into the optimization objective more explicitly. The present benchmark uses effort and saturation terms, but tendon rate, fatigue, friction, backlash, routing hysteresis, and motor current are not modeled. For hardware deployment, these variables can be as important as curvature error. A third direction is to increase the state fidelity from piecewise constant strain to higher-order strain bases or spatial discretization while retaining a reduced controller state.

4.6 Limitations and future work

The value of the reported benchmark depends on stating its limits clearly. The present work is a numerical control study, not an experimental validation, and several assumptions should be addressed before drawing hardware-level conclusions.

First, the supplied optimizer results use one stochastic run per method. The random seed is fixed for reproducibility, but one run cannot support inferential claims about the probability that GWO, POA, or another stochastic algorithm is universally superior. A journal extension focused on optimizer statistics should use at least 20–30 independent runs, report median and interquartile range, and apply nonparametric paired tests with multiple-comparison correction. The current study instead uses the eleven optimizers as alternative gain-generation mechanisms within a fidelity benchmark.

Second, the physical parameters are representative model values rather than identified measurements from a specific hardware specimen. The effective modulus, distributed mass, disk mass, tendon routing radius, preload, and payload should be measured for experimental validation. Material viscoelasticity, tendon friction, sheath friction, backlash, motor dynamics, tendon elasticity, and cable routing losses are omitted. Their addition will generally increase the discrepancy between ideal inverse dynamics and hardware response.

Third, the Cosserat reduction assumes shear-inextensible, torsion-free piecewise constant strain within each section. This is appropriate for a compact benchmark but does not represent arbitrary distributed strain. Large localized loads or contact could require a higher-order strain basis, finite-element Cosserat discretization, or spatial shooting/collocation method.

Fourth, the full-model fidelity audit is intentionally severe because the controller feedforward is still constructed from the reference lookup while the simulated plant recomputes nonlinear dynamics at the actual state. The resulting 400.589 mm tip RMSE should therefore be read as evidence that the current feedforward is not self-consistent under large departure from the reference. Future work should close this loop by recomputing inverse dynamics online or by using a state-dependent surrogate. The large gap is not claimed as an intrinsic limitation of Cosserat control; it is a quantified limitation of the present reduced lookup strategy.

Fifth, the disturbance cases are synthetic generalized loads and measurement noise. They provide controlled stress tests but are not replacements for experimentally measured contact forces, payload uncertainty, or sensor error distributions. Hardware studies should use force/torque sensing or calibrated payloads and report the actual disturbance spectrum.

Future work will therefore proceed in four stages: (1) repeated-run statistical optimizer evaluation; (2) state-dependent or multi-fidelity inverse-dynamics compensation; (3) inclusion of tendon/motor dynamics and friction; and (4) experimental validation on a two-section cable-driven prototype with synchronized shape sensing and tendon-force measurement. These steps directly follow from the failure modes exposed by the present benchmark.

5 Conclusions

This paper presented a fidelity-aware inverse-dynamics and PID-control benchmark for a two-section tendon-driven continuum robot described by a reduced shear-inextensible Cosserat model. Six tendon tensions were mapped to four generalized curvature coordinates under preload and saturation limits, while inverse-dynamics feedforward was combined with section-wise PID feedback. Eleven established metaheuristics were compared under identical search conditions, and all tuned controllers were cross-verified using RK4 and Newmark-β integration and four disturbance scenarios.

The nominal RK4 benchmark identified GWO and POA as the strongest observed controllers, with tip RMSE values of 24.699 and 24.945 mm, respectively. Their advantage remained under payload, noise, and combined disturbances, with POA reaching 25.062 mm and GWO 27.265

mm in the combined condition. The study also showed that terminal tuning cost alone is an unreliable proxy for final Cartesian accuracy: ACO obtained an objective value almost identical to GWO while producing more than six times the RK4 tip RMSE. Numerical cross-verification exposed even stronger sensitivity for GA and BOA, whose Newmark results were approximately 16–17 mm despite RK4 errors near 155–156 mm.

The most important result emerged from the full configuration-dependent Cosserat replay. The GWO controller's tip RMSE increased from 24.699 mm on the reference-linearized model to 400.589 mm when mass, gravity, and Coriolis quantities were recomputed at the actual configuration. This finding changes the role of the benchmark from a simple optimizer competition to a fidelity audit. It demonstrates that successful tuning on a computationally convenient model must be followed by independent numerical and nonlinear-model verification before strong control claims are made.

Accordingly, the principal contribution of the work is an integrated and reproducible validation chain rather than a new optimizer. The framework provides a common basis for future comparisons of adaptive control, MPC, learning-based methods, and multi-fidelity optimization while retaining physically bounded tendon actuation and explicit checks on numerical sensitivity.

List of abbreviations

ABC	Artificial bee colony
ACO	Ant colony optimization
BOA	Butterfly optimization algorithm
CS	Cuckoo search
GA	Genetic algorithm
GWO	Grey wolf optimizer
HHO	Harris hawks optimization
ITAE	Integral of time-weighted absolute error

MSE	Mean squared error
PID	Proportional–integral–derivative
POA	Pelican optimization algorithm
PSO	Particle swarm optimization
RK4	Fourth-order Runge–Kutta method
RMSE	Root-mean-square error
TLBO	Teaching-learning-based optimization
WOA	Whale optimization algorithm

Declarations

Availability of data and materials

The datasets and MATLAB code used and/or analysed during the current study are available from the corresponding author on reasonable request.

Competing interests

The authors declare that they have no competing interests.

Funding

Not applicable.

Authors' contributions

Both authors jointly wrote the manuscript. All authors read and approved the final manuscript.

Acknowledgements

Not applicable.

References

[1] Russo M, Sadati SMH, Dong X, Mohammad A, Walker ID, Bergeles C, Xu K, Axinte DA (2023) Continuum Robots: An Overview. Advanced Intelligent Systems 5(5):2200367. https://doi.org/10.1002/aisy.202200367

[2] Alandoli EA, Fan Y, Liu D (2025) A review of extensible continuum robots: mechanical structure, actuation methods, stiffness variability, and control methods. Robotica 43(2):764-791. https://doi.org/10.1017/S0263574724001966

[3] Fan Y, Yi B, Liu D (2024) An overview of stiffening approaches for continuum robots. Robotics and Computer-Integrated Manufacturing 90:102811. https://doi.org/10.1016/j.rcim.2024.102811

[4] Sincak PJ, Prada E, Mikova L, Varga M, Merva T, Virgala I (2024) Sensing of Continuum Robots: A Review. Sensors 24(4):1311. https://doi.org/10.3390/s24041311

[5] Djeffal S, Mahfoudi C (2023) Inverse kinematic model of multi-section continuum robots using particle swarm optimization and comparison to four meta-heuristic approaches. SIMULATION 99(8):817-830. https://doi.org/10.1177/00375497231164645

[6] Djeffal S, Ghoul A, Saadi A, Izri Z, Morakchi MR, Wang H (2025) Design optimization and simulation of a 3D printed cable-driven continuum robot using IKM-ANN and nTop software. Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering 239(2):200-218. https://doi.org/10.1177/09596518241276233

[7] Djeffal S, Morakchi MR, Ghoul A, Kargin TC (2024) DDPG-based reinforcement learning for controlling a spatial three-section continuum robot. Franklin Open 6:100077. https://doi.org/10.1016/j.fraope.2024.100077

[8] Djeffal S, Mahfoudi C, Ghoul A (2024) Solving the Non-Linear Dynamic Equations of Motion of a Variable Curvature Continuum Robot Using Runge-Kutta Method and MATLAB Software. Journal of Applied Nonlinear Dynamics 13(1):13-26. https://doi.org/10.5890/JAND.2024.03.002

[9] Aner EA, Awad MI, Shehata OM (2024) Performance evaluation of PSO-PID and PSO-FLC for continuum robot's developed modeling and control. Scientific Reports 14:733. https://doi.org/10.1038/s41598-023-50551-0

[10] Ba W, Chang J, Liu J, Wang X, Dong X, Axinte D (2024) An analytical differential kinematics-based method for controlling tendon-driven continuum robots. Robotics and Autonomous Systems 171:104562. https://doi.org/10.1016/j.robot.2023.104562

[11] He S, Zou C, Deng Z, Liu W, He B, Zhang J (2024) Model-less optimal visual control of tendon-driven continuum robots using recurrent neural network-based neurodynamic optimization. Robotics and Autonomous Systems 182:104811. https://doi.org/10.1016/j.robot.2024.104811

[12] Feizi N, Pedrosa FC, Jayender J, Patel RV (2025) Deep Koopman Approach for Nonlinear Dynamics and Control of Tendon-Driven Continuum Robots. IEEE Robotics and Automation Letters 10(3):2926-2933. https://doi.org/10.1109/LRA.2025.3533967

[13] Alora JI, Cenedese M, Haller G, Pavone M (2025) Discovering dominant dynamics for nonlinear continuum robot control. npj Robotics 3:5. https://doi.org/10.1038/s44182-025-00021-8

[14] Danesh R, Janabi-Sharifi F (2025) Backstepping control of tendon-driven continuum robots in large deflections using the Cosserat rod model. Mechanism and Machine Theory 208:105953. https://doi.org/10.1016/j.mechmachtheory.2025.105953

[15] Liu H, Liu W, Zhong J (2026) Adaptive neural network fault-tolerant control of tendon actuated continuum robots under actuator failures and external disturbances. Robotics and Autonomous Systems 198:105310. https://doi.org/10.1016/j.robot.2025.105310

[16] Maghooli N, Mahdizadeh O, Bajelani M, Moosavian SAA (2025) Learning-based control for tendon-driven continuum robotic arms. Frontiers in Robotics and AI 12:1488869. https://doi.org/10.3389/frobt.2025.1488869

[17] Huertas Nino MP, Boutayeb M, Martinez D (2025) A hybrid tendon-driven continuum robot that avoids torsion under external load. Frontiers in Robotics and AI 12:1576209. https://doi.org/10.3389/frobt.2025.1576209

[18] Yang W, Liu Z, Cao Y, Wang S, Xie L (2025) Design and control of a continuum robot with switchable stiffness based on ball-and-socket joints. Scientific Reports 15:34363. https://doi.org/10.1038/s41598-025-16987-2

[19] Ouyoucef A, Peyron Q, Zheng G, Boyer F (2025) Durability-Aware Trajectory Planning and Quasi-Static Control of a Continuum Parallel Robot for Industrial Applications. IEEE Robotics and Automation Letters 10(11):12016-12023. https://doi.org/10.1109/LRA.2025.3619707

[20] Ghoul A, Sattar AM, Adoul MA, Sadki O, Djeffal S (2026) Obstacle-aware inverse kinematics of variable-length continuum robots via teaching-learning-based optimization with experimental validation. Scientific Reports 16:11227. https://doi.org/10.1038/s41598-026-46132-6

[21] Jabari M, Visconte C, Quaglia G, Laribi MA (2026) Hybrid Supervised Learning and Constrained Optimization for Task-Space Trajectory Generation in Tendon-Driven Continuum Robots. ASME Open Journal of Engineering 5:051003. https://doi.org/10.1115/1.4070630

[22] Arefinia E, Feizi N, Pedrosa FC, Patel RV, Jayender J (2026) Error-State Model Predictive Path Integral Control of Tendon-Driven Continuum Robots Using Cosserat Rod Dynamics With Strain Parametrization. IEEE Robotics and Automation Letters 11(3):3867-3874. https://doi.org/10.1109/LRA.2026.3662658

[23] Jabari M, Visconte C, Quaglia G, Laribi MA (2026) Prescribed-time fault-tolerant optimal control of continuum surgical robots via composite deep learning. European Journal of Control 90:101514. https://doi.org/10.1016/j.ejcon.2026.101514

[24] Djeffal S, Morakchi MR, Bendriss B, Talhi Y, Ghoul A (2026) Machine learning for precise continuum robot motion with experimental validation. Franklin Open 16:100670. https://doi.org/10.1016/j.fraope.2026.100670