


## RESEARCH ARTICLE OPEN ACCESS

# Generalized Task-Driven Design of Soft Robots via Reduced-Order Finite Element Method-Based Surrogate Modeling

Yao Yao<sup>1</sup> | David Howard<sup>2</sup> | Perla Maiolino<sup>1</sup> <sup>1</sup>Oxford Robotics Institute, University of Oxford, Oxford, UK | <sup>2</sup>CSIRO Robotics, CSIRO, Pullenvale, Queensland, Australia**Correspondence:** Perla Maiolino ([perla.maiolino@eng.ox.ac.uk](mailto:perla.maiolino@eng.ox.ac.uk))**Received:** 21 January 2026 | **Revised:** 27 March 2026 | **Accepted:** 1 April 2026**Keywords:** design | modeling | optimization | task-driven design

## ABSTRACT

Task-driven design of soft robots requires models that are physically accurate and computationally efficient, while remaining transferable across actuator designs and task scenarios. However, existing modeling approaches typically face a fundamental trade-off between physical fidelity and computational efficiency, which limits model reuse across design and task variations and constrains scalable task-driven optimization. This paper presents a unified reduced-order finite element method (FEM)-based surrogate modeling pipeline for generalized task-driven soft robot design. High-fidelity FEM simulations characterize actuator behavior at the modular level, from which compact surrogate joint models are constructed for evaluation within a pseudo-rigid body model (PRBM). A meta-model maps actuator design parameters to surrogate representations, enabling rapid instantiation across a parameterized actuator family. The resulting models are embedded into a PRBM-based simulation environment, supporting task-level simulation and optimization under realistic physical constraints. The proposed pipeline is validated through sim-to-real transfer across multiple actuator types, including bellow-type pneumatic actuators and a tendon-driven soft finger, as well as two task-driven design studies: soft gripper co-design via reinforcement learning and 3D actuator shape matching via evolutionary optimization. The results demonstrate high accuracy, efficiency, and reliable reuse, providing a scalable foundation for autonomous task-driven soft robot design.

## 1 | Introduction

Soft robots are increasingly valued for their adaptability, compliance, and ability to achieve a wide range of behaviors through the combination of material properties, morphology, and control [1–3]. However, these advantages also make the design of soft robots a significant challenge, leading to vast and complex design spaces where manual or heuristic approaches are typically insufficient to ensure optimal performance in many task applications. This has led to an increasing focus on task-driven design, wherein morphology and control are systematically tailored to meet specific application objectives via computational approaches [4, 5].

Such approaches rely heavily on accurate and scalable modeling to explore large design spaces. However, soft systems pose

particular difficulties in this context, especially when predicting forces and self or environmental interactions, due to their inherent softness and material nonlinearity. This not only challenges the commonly discussed issue of accuracy and efficiency, but also reveals a critical modeling bottleneck: the lack of generalizable models that can be reused across parameterized geometries, actuation types, and task scenarios. Here, generalization refers to the ability of a model to retain accuracy across variations within a parameterized design space and across different task-level conditions represented in simulation, rather than being confined to case-specific configurations.

This limitation stems from a fundamental trade-off in modeling soft robots. Despite a wide variety of modeling techniques having

---

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Advanced Intelligent Systems* published by Wiley-VCH GmbH.

been explored, including first-principles-based methods, physics-based simulation, and data-driven or reduced-order models, most approaches follow one of two dominant directions: high-fidelity modeling and simplified or reduced-order representations. Approaches based on high-fidelity modeling, such as finite element method (FEM), offer accurate predictions but are often too computationally expensive for repeated task-level evaluation, which in practice encourages the use of simplified or proxy objectives tied to specific tasks and assumptions [6, 7]. In contrast, simplified or reduced-order models are computationally efficient, but typically rely on assumptions and approximations for particular designs or operating conditions, limiting their transferability to new geometries, loading regimes, or task settings [8, 9].

As a result, existing approaches fall short of providing models that are simultaneously accurate and efficient while remaining sufficiently general for task-driven design. This lack of generalization limits model scalability and deployability across broader design contexts, thereby hindering the development of general-purpose modeling tools that the research community can readily build upon. Addressing this limitation is essential for building flexible and reusable design pipelines in soft robotics.

To overcome this challenge, this paper introduces a unified modeling pipeline for task-driven design of soft robots. Specifically, we:

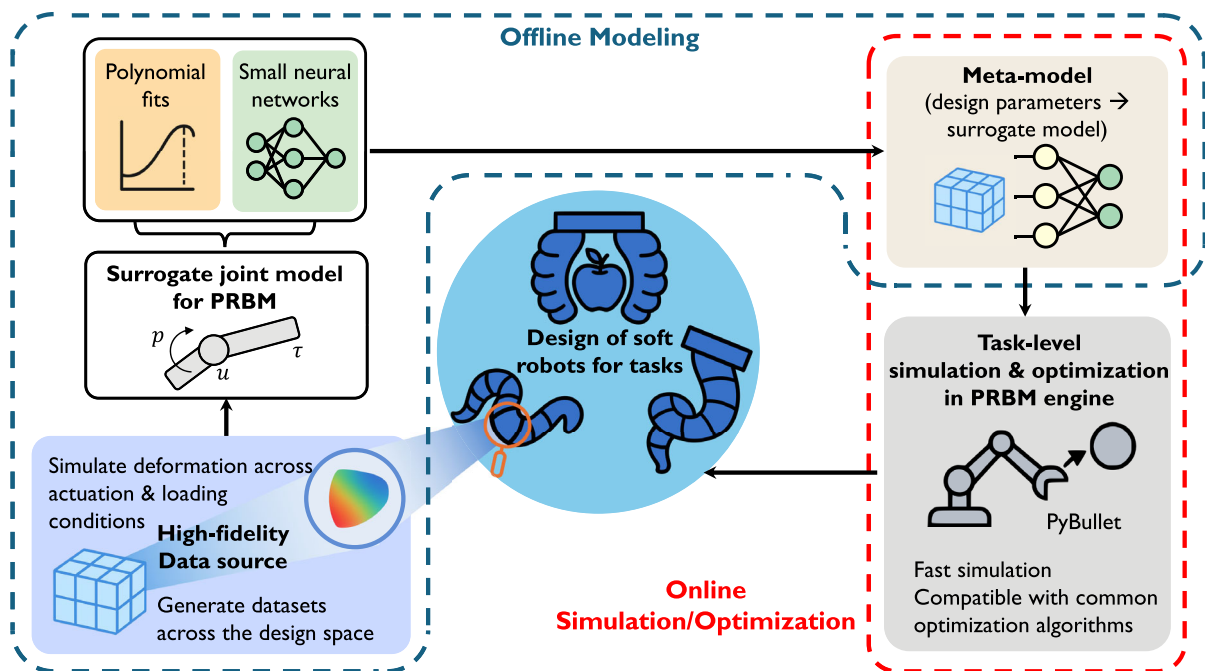
- Use high-fidelity FEM simulations to generate actuator data that captures coupled relationships between actuation, deformation, and force at the modular level;
- Create surrogate models from this data using either polynomial regressions or small (NN)s, serving as joint models within the pseudo-rigid body model (PRBM) framework, to

significantly accelerate simulation while maintaining FEM-level accuracy;

- Develop a meta-model that maps actuator design parameters to surrogate features, enabling rapid model interpolation within a parameterized actuator family;
- Integrate the resulting models into a fast and flexible PRBM-based simulation environment, such as PyBullet, supporting efficient task-level simulation and optimization under real-world physical constraints such as contact, loading, and gravity.

Together, these contributions support pipeline reuse across design variations and task conditions, offering a practical pathway toward scalable task-driven design. Notably, the proposed approach is scoped to parameterized actuator families and is not intended to generalize across fundamentally different soft robotic morphologies or actuation principles.

We illustrate our contributions using a range of different soft actuators, including pneumatically actuated bellow-type soft actuators, which are widely used in the soft robotics field due to their large deformation, adaptability, straightforward control, and 3D printing compatibility. These actuators remain challenging to model accurately while efficiently due to effectively infinite degrees of freedom and distributed actuation [10]. We also include a tendon-driven finger, demonstrating that the proposed pipeline extends beyond a single actuator type. Across diverse applications, including soft gripping and 3D shape-matching, we demonstrate that the pipeline supports both accurate physical modeling and efficient task-level optimization. An overview of the workflow is provided in Figure 1.



**FIGURE 1** | Schematic overview of the proposed pipeline for modeling and task-driven design of soft robots. The process starts from a well-parameterized modular actuator design, which is simulated using FEM to generate high-fidelity training data. Compact surrogate models are then constructed from these data as joint models, represented using either polynomial fits or small NNs, for use in PRBM-based simulations. By repeating this process across different designs, meta-models are constructed to capture the mapping from design parameters to surrogate models. The resulting surrogate and meta-models can be integrated with PRBM-based simulators and state-of-the-art (SOTA) optimization algorithms, enabling fast task-level simulation and task-driven design optimization.

## 2 | Literature Review

Modeling plays a critical role in the design of soft robots, as it directly impacts both the accuracy of the predicted behavior and the speed of the automated design process. A wide range of modeling approaches is available, each with its own trade-offs between accuracy, efficiency, and generalizability.

Analytical models, such as rod formulations [11] or the minimum potential energy principle [12], are valued for their simplicity and computational efficiency. However, they often rely on simplifying assumptions, such as uniform curvature, constant cross-sections, or linearized material behavior. These assumptions limit flexibility and often require re-derivation or recalibration when designs, materials, or boundary conditions change.

Geometric models, such as piecewise constant curvature and function-based curves, provide efficient representations of the kinematics of soft actuators and intuitive mappings from actuation inputs to robot configurations [13]. However, they tend to break down under unexpected external forces such as gravity or contact [9].

Discrete models, such as PRBMs, approximate continuum deformation by connecting rigid segments with compliant joints [14]. They are lightweight and compatible with existing rigid-body simulators (e.g., PyBullet), enabling reliable and efficient task- and environment-level simulation [15, 16]. However, they typically require careful parameter calibration to map the lumped behavior back to distributed soft structures. This process is time-consuming and often lacks generalizability beyond the specific robot for which the model is tuned. Such limitations highlight the need for modeling approaches that can be reused across actuator designs, enabling broader applicability within the research community.

Numerical methods, especially FEM, offer the highest fidelity by solving governing equations over the continuum. FEM is widely used for simulating soft robots because it provides high accuracy and strong generality, accommodating arbitrary geometries, materials, actuation, and boundary conditions, with established tool support such as COMSOL Multiphysics, ANSYS, and Abaqus [7, 17, 18]. Other numerical approaches have also been explored for deformable simulation, particularly in the computer graphics community, including extended position-based dynamics (XPBD) [19] and material point methods (MPM) [20]. XPBD enables efficient and stable simulations but often yields less reliable mechanical responses, while MPM can capture large deformations but may introduce numerical diffusion that affects accurate force responses. However, the high computational cost and convergence challenges of FEM under complex contact and large-deformation conditions limit its applicability in large-scale design optimization. As a result, task-level performance is often approximated through actuator-level objectives that are more tractable to evaluate but are closely tied to specific modeling assumptions and task formulations. This limitation has motivated efforts in model order reduction (MOR) techniques, which seek to retain fidelity while reducing computational expense [6, 21].

MOR techniques aim to accelerate FEM by reducing system dimensionality while retaining dominant deformation modes. Examples include proper orthogonal decomposition [8, 22], condensed FEM formulations [9], and Cosserat rod implementations

such as SoRoSim [23]. These methods improve computational efficiency and can reproduce complex deformations with fewer degrees of freedom. However, the reduced bases or governing formulations are typically derived for specific geometries, boundary conditions, or actuation schemes, making the resulting models configuration-specific and difficult to reuse in task-driven design workflows.

Surrogate models learn input–output relationships from high-fidelity simulations or experiments, providing fast evaluation once trained. Approaches include regression models, Gaussian processes, and NNs that map actuation to deformation and force responses [24, 25]. Compared with physics-based reductions, data-driven surrogates can offer greater flexibility. However, their performance is inherently tied to the coverage of the training data, trained under fixed loading conditions or tailored to a specific design, making it hard to extrapolate reliably when tasks, geometries, or operating regimes change.

In summary, existing approaches present clear trade-offs. Analytical, geometric, and discrete models are efficient but rely on restrictive assumptions that limit accuracy and transferability. FEM provides accuracy and generality across arbitrary geometries and boundary conditions, but is too costly to support large-scale task-level design optimization. Physics-based model reduction improves efficiency but remains tied to specific designs or loading conditions. Data-driven surrogates accelerate evaluation and can capture complex behaviors, yet they are often restricted to fixed training scenarios and lack generalization across tasks. Overall, what is still missing is a pipeline that unifies these ideas into models that are accurate, efficient, and reusable across designs and task scenarios. Addressing this gap is essential not only for advancing task-driven design of soft robots but also for providing a modeling foundation that can be extended and adopted by the broader soft robotics community.

## 3 | Materials and Methods

This section introduces the proposed modeling pipeline for task-driven design of soft robots. The pipeline combines four components: (i) high-fidelity FEM simulations to capture coupled actuation-deformation-force responses; (ii) surrogate models that encode these behaviors into compact representations for fast evaluation, and (iii) PRBM-based simulation in PyBullet, embedding surrogates for task-level simulation. (iv) meta-models that map design parameters to surrogate models, supporting scalability across new actuator designs.

To explicitly illustrate this workflow, as presented in Figure 1, the pipeline starts from a parameterized modular actuator design, whose behavior is first characterized using high-fidelity FEM under both free-loading and constrained conditions. Building on the FEM data, compact surrogate joint models are derived to approximate the nonlinear relationship between actuation input and mechanical response at the modular level for use within a PRBM-based simulation. These surrogates significantly reduce computational cost while preserving physical fidelity.

To support design generalization, the surrogate construction process is repeated across multiple actuator designs, and the resulting model parameters are learned by a higher-level meta-model that maps from design parameters to surrogate representations.

This enables rapid on-demand generation of behavior models for unseen designs via interpolation within a parameterized actuator family, eliminating the need for repeated FEM simulations. Finally, the design-conditioned surrogate models are embedded into the PRBM-based simulation in PyBullet, enabling task-level simulation with realistic physical interactions. This integrated environment also supports efficient task-level design using SOTA learning and optimization algorithms.

### 3.1 | FEM Data Generation

Existing surrogate models are often developed for specific loading conditions, such as free deformation, and may not capture the coupled effects of actuation and external forces. However, task-level simulations within a PRBM framework require joint laws that reflect both deformation and the corresponding forces. To address this, we design FEM simulations that provide a comprehensive mapping between actuation input, deformation, and reaction forces at the modular level.

The FEM simulations are conducted under quasi-static assumptions to characterize the dominant elastic response of the actuator modules. This is appropriate for the considered systems, as the actuation and deformation processes occur at relatively low speeds, and inertial effects are negligible compared to elastic forces. The resulting actuation-deformation-force relationships are used as joint models, which are subsequently embedded into task-level simulations to account for gravity, contact, and external interactions.

The procedure consists of two complementary conditions. In the free-loading condition, actuation inputs are incrementally applied to the actuator while external loads are absent, yielding the natural actuation-deformation response. In the constrained condition, the actuation input is held fixed while external forces are applied to vary the deformation, enabling explicit capture of the force-deformation relationship under constraint. The FEM data are sampled over actuation ranges determined by the intended task operating conditions, while the external force range is constrained by stable FEM convergence. These simulations collectively yield coupled actuation-deformation-force data used to derive compact surrogate models.

For clarity, an illustrative example using a translational bellow-soft pneumatic actuator (SPA) is provided in the Supporting Information 1.1, where specific FEM configurations, parameter sweeps, and numerical results are detailed.

### 3.2 | Surrogate Model via Polynomial Fits

The FEM data described above provides actuator states under both free-loading and constrained conditions. To enable efficient task-level simulation within a PRBM framework, we derive a compact analytical surrogate that represents the joint force as a function of actuation and deformation. Here  $\tau$  denotes the generalized joint effort associated with the deformation coordinate  $u$ . Depending on the actuator configuration, this effort corresponds to a torque for rotational deformation or a force for translational deformation. This surrogate is constructed via polynomial approximations of stiffness terms obtained from FEM data.

Based on the principle of virtual work, the joint force  $\tau$  can be expressed as the derivative of the actuator's potential energy  $U$  with respect to the deformation  $u$ , which itself depends on both the actuation parameter  $p$  and the deformation  $u$ :

$$\tau(p, u) = \frac{dU(p, u)}{du} = \frac{\partial U(p, u)}{\partial u} + \frac{\partial U(p, u)}{\partial p} \frac{dp(u)}{du} \quad (1)$$

Given the quasi-static FEM results, this derivative can be approximated using finite differences between adjacent simulation states. Distinguishing between free-loading (actuation varied, no external force) and constrained conditions (actuation fixed, external force applied), the joint forces can be decomposed into actuation and elastic contributions. By differentiating further with respect to  $u$ , we obtain stiffness terms  $k_a(p)$  and  $k_e(u)$ , which depend solely on actuation and deformation, respectively. This separability allows each term to be independently approximated using low-order polynomial fits. Without this decomposition, the joint force would need to be approximated using a single high-order polynomial over the combined actuation-deformation space. Such models typically require significantly more data to avoid overfitting and may lead to unstable behavior when embedded within iterative physics simulations. By separating the actuation-dependent and deformation-dependent stiffness terms, the surrogate fitting problem becomes lower-dimensional and more data-efficient. Such a decomposition constitutes an approximation, as it neglects higher-order coupling between actuation and deformation, and is expected to remain valid within the moderate pressure and deformation ranges considered in this study.

The resulting surrogate provides an explicit, closed-form mapping from  $(p, u)$  to the net force:

$$\tau_n(p, u) = - \int_0^u (k_a(p) + k_e(u)) du + \tau_n(p, 0), \quad (2)$$

where  $\tau_n(p, 0)$  is determined from free-loading conditions. All functions are represented by polynomial regressions fitted to FEM data, yielding a lightweight model suitable for use in rigid-body-based physics engines. Detailed derivations and an illustrative case study on a translational bellow-SPA are provided in the Supporting Information 1.2 and 1.3. When extending the module-level surrogate to multi-module actuators through module-wise composition, additional system-level limitations arise. In particular, cumulative inter-module effects—such as compliance accumulation and boundary constraints—are not explicitly modeled, leading to a systematic underestimation of deformation as the number of modules increases, as quantified in Supporting Information Figure S2f.

### 3.3 | Surrogate Model via Small Neural Networks

While polynomial fits offer interpretable closed-form surrogates, their approximation accuracy is inherently limited, especially when actuator behavior becomes strongly nonlinear. For example, in bellow-SPA modules with internal constraints, even higher-order fits struggle to capture the underlying relationships accurately (as demonstrated via Supporting Information Figure S3). To address this, we employ small NNs as alternative surrogates.

The network takes the actuation parameter  $p$  and deformation  $u$  as inputs and outputs the corresponding net joint force  $\tau_n(p, u)$ :

$$(p, u) \rightarrow \tau_n(p, u) \quad (3)$$

Unlike polynomial regression, NNs flexibly approximate nonlinear relationships without requiring explicit functional assumptions. Importantly, the networks are designed to remain lightweight, ensuring efficient evaluation within PRBM-based task simulations.


After training, the NNs are directly integrated into rigid-body-based physics engines such as PyBullet for task simulation, where they are used as interpolation models within the sampled operating ranges of actuation and deformation, extending applicability to more complex actuator responses.

We implement small feed-forward NNs trained on FEM-derived data spanning both free-loading and constrained conditions. They provide significantly improved accuracy compared with polynomial fits, while remaining lightweight for deployment in task simulations. However, compared with the analytical polynomial surrogate, NN-based models may produce less predictable responses when evaluated slightly outside the training distribution in physics-based simulations, and incur higher computational cost during evaluation. Detailed architecture, training setup, and validation on a representative bellow-SPA are provided in the Supporting Information 1.4.

### 3.4 | Meta-Model for Design-Conditioned Surrogate Generation

The purpose of the meta-model is to establish a direct mapping from actuator design parameters (e.g., geometric variables such as module radius, length, or thickness) to the corresponding reduced-order surrogate representations, enabling task-level simulations for previously unseen designs through design-space interpolation without repeating high-fidelity FEM simulations. By learning this design-to-behavior relationship, the meta-model enables rapid generation of design-conditioned surrogate models and therefore serves as the key interface between offline physical modeling and online task-driven design and optimization.

For surrogate models constructed using polynomial fits, the meta-model is implemented as a neural network that maps the actuator design parameters to the corresponding set of polynomial coefficients, as illustrated in Figure 2a. In this way, the



**FIGURE 2** | Meta-models for design-conditioned surrogate generation. (a) A NN mapping the design parameters to the corresponding polynomial coefficients of the surrogate model. (b) A NN mapping the combined design parameters and modular actuation inputs to the resulting net force.

polynomial surrogate for a new design can be generated instantly from its design parameters, allowing fast evaluation of actuator behavior within the simulation environment.

For surrogate models constructed using small NNs, all available data across different designs are directly aggregated to form a unified training dataset. For each sample, the design parameters and the modular variables  $(p, u)$  are concatenated as the inputs, and the corresponding net force  $\tau$  is taken as the output. A NN is then trained to capture the global relationship between the combined design-actuation inputs and the resulting behavioral response, as shown in Figure 2b. This formulation enables the surrogate to generalize jointly over both the design space and the actuation space within a single compact model.

### 3.5 | Use of AI Tools

The web version of ChatGPT (OpenAI; model: GPT-5) was used solely for English language editing and improving clarity of phrasing. The authors reviewed and edited all AI-assisted text and take full responsibility for the final content.

## 4 | Results and Discussion

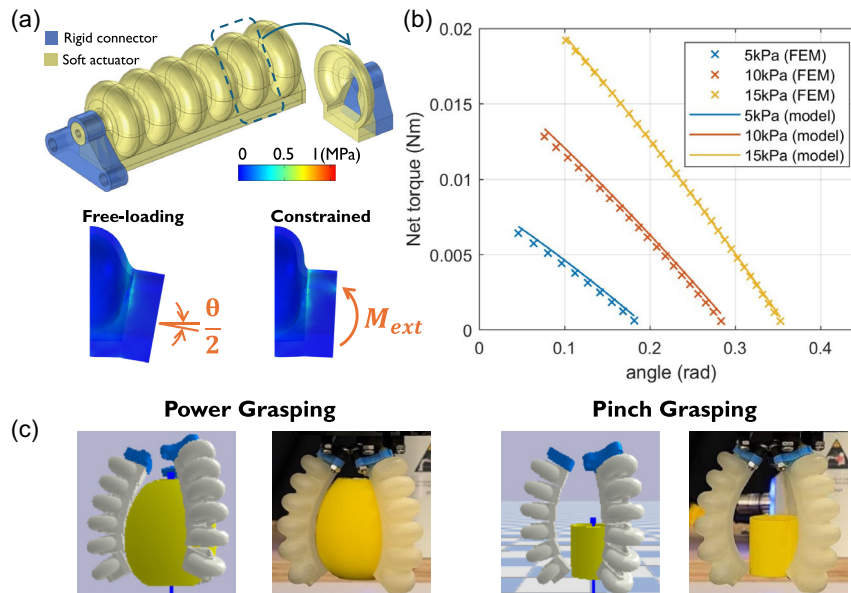
This section reports the performance of the proposed approach. Two complementary evaluations are presented: (i) the sim-to-real transfer of the surrogates on tasks, and (ii) their ability to support task-driven design through optimization. The goal of these evaluations is to demonstrate that the proposed pipeline is not tailored to a single application, but instead offers a general modeling and optimization approach capable of supporting heterogeneous task requirements. This is addressed by applying the pipeline to tasks that differ fundamentally in both physical interaction and optimization objectives. Specifically, grasping is a contact-based task in which forces are mainly transmitted through friction, and success is defined by the stability of the grasp and the ability to lift objects under load, whereas 3D shape-matching requires continuous geometric accuracy under distributed loading and gravity. The following subsections present detailed, case-specific analyses for each task and actuator type. A brief general discussion at the end of this section summarizes shared assumptions and limitations across these evaluations.

### 4.1 | Surrogate Models in Task Simulations

This subsection focuses on the sim-to-real validation of the proposed surrogate models in several task simulations involving complex deformations and contact interactions. The objective is to verify that the surrogate models can accurately reproduce actuator behavior across different actuation principles and task types, without relying on full FEM simulations at runtime.

#### 4.1.1 | Contact-Based Grasping With Polynomial Surrogate

To evaluate the capability of the surrogate models constructed using polynomial fits in task-level simulation, and to validate the proposed pipeline on a representative class of contact-dominated manipulation tasks, we simulated a soft gripper performing grasping, where the smooth and efficient numerical



**FIGURE 3** | Validation of the polynomial-based surrogate model through a soft gripper grasping task. (a) FEM model of the bellow-SPA used in the gripper. (b) Comparison of net torque versus angular deflection under different pressure levels obtained from the FEM simulation and the surrogate model. (c) Two grasping modes with different objects are compared between simulation and experiment.

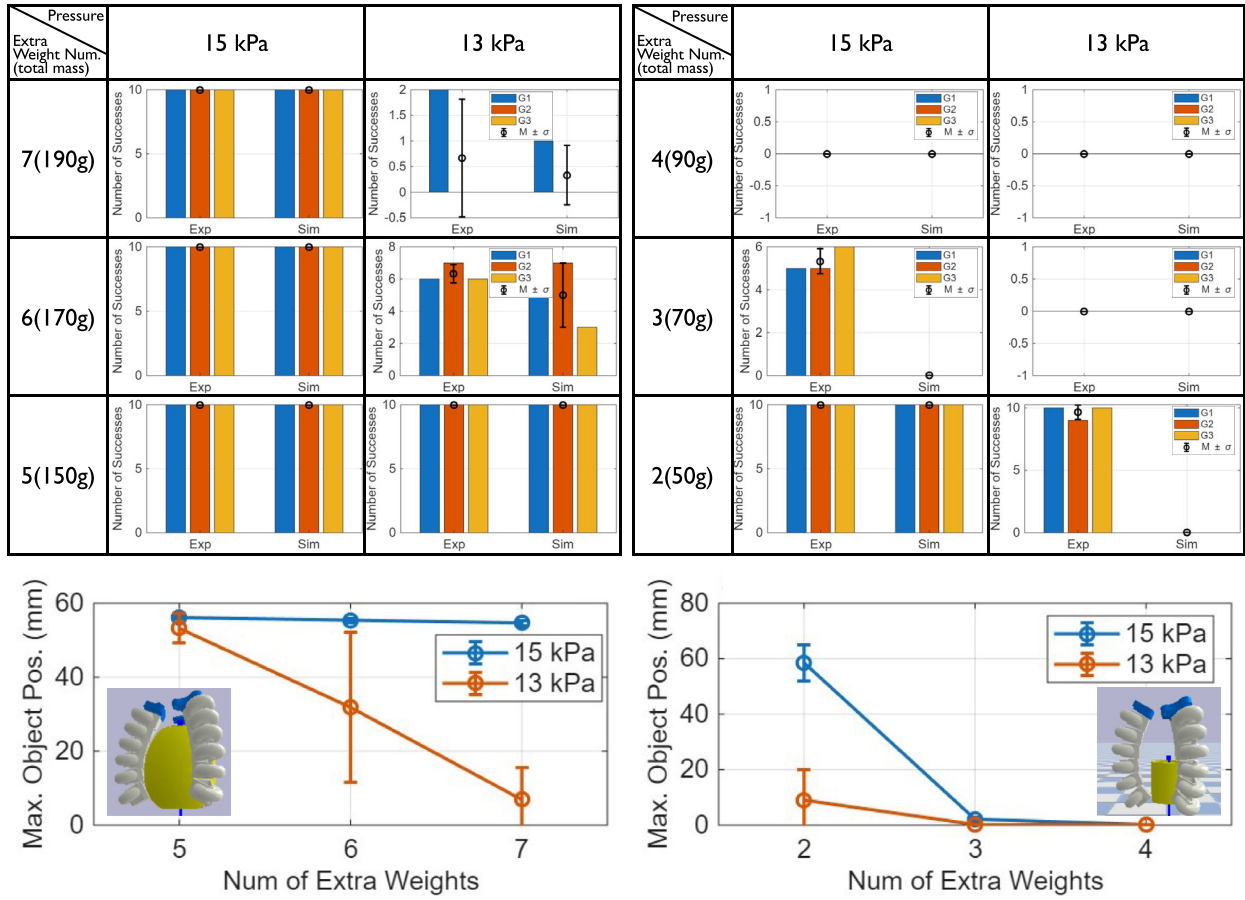
characteristics of polynomial surrogates facilitate stable physics-based simulation. The gripper design is based on the default bellow-SPA from prior work, featuring a strain-limiting side that constrains expansion on one face and thereby promotes directional bending during pressurization [26].

As illustrated in Figure 3a, we simulated a half-module, closed with an end cap to facilitate smooth deformation, under conditions of free pressurization and externally applied moment  $M_{ext}$  at various fixed internal pressures, obtaining angular deflection  $\theta$ . In this study, the pressure is sampled from 0 to 15 kPa with 0.5 kPa as an interval, and the externally applied moment is sampled from 0 to 0.02 N · m with an interval of 0.0006 N · m. Following the previously described modeling process, we compared the model-predicted net torque to the applied external moment across different pressure values, resulting in a maximum error of 8.26% relative to the moment range.

Grasping simulations were then carried out in PyBullet using a physics-based model of the soft gripper. The gripper consists of three actuator fingers equally spaced around a central axis, with each finger modeled as a chain of links connected by revolute joints. Visual and collision geometry are taken directly from the actuator CAD STL mesh files, while the mass properties of each link are estimated from the actuator volume and material density. The lateral friction coefficient between the actuator and PLA-printed objects is set to 0.8 [27], and anchor friction is enabled to mimic soft contact behavior. To better approximate real-world conditions, 5% random noise is introduced to the object placement, friction coefficient, and actuation pressure in the simulations. To ensure a fair sim-to-real comparison, the simulated grasping setup is configured to match the physical experiments in terms of gripper geometry, object placement, and actuation conditions. Details of the ablation study used for selecting the simulation parameters, together with the corresponding single-run simulation time, are reported in the Supporting Information 2.1.

Using this simulation setup, we evaluated sim-to-real transfer by comparing grasping performance against physical experiments. Two representative objects were considered: a 3D-printed egg-shaped object (maximum radius: 30 mm; height: 65 mm) for power grasping, and a cylindrical object (radius: 16 mm; height: 35 mm) for pinch grasping. The objects weighed 50 and 10 g, respectively. By systematically increasing the object mass by adding standardized 20 g weights (at a predefined central location integrated into the object design prior to fabrication) and applying different actuation pressures, we evaluated how closely the simulated grasping success rates align with experimental outcomes. Specifically, grasping was conducted under two actuation pressures, 15 and 13 kPa. For each grasping mode, three weight levels were tested: for power grasping of the egg-shaped object, 7, 6, and 5 units of 20 g were added, while for pinch grasping of the cylinder, 4, 3, and 2 units of 20 g were used. Despite the small sample sizes and the use of binary and coarse success metrics, this sim-to-real validation is intended to demonstrate the feasibility of transferring surrogate-based task-level simulation outcomes to physical grasping behavior, rather than to provide a comprehensive performance benchmark.

Physical experiments were conducted using a Franka Emika Panda arm (Franka Robotics, Germany) with the soft gripper mounted as the end effector and controlled through its standard interface. The actuation pressure was supplied by an Anest Iwata air compressor (Anest Iwata Corporation, Japan) and regulated by a VPPE-3-1-1/8-2-010-E1 Festo proportional-pressure regulator (Festo AG & Co. KG, Germany). At each combination of weight level and pressure setting, three sets of ten grasping trials were performed, yielding three success rates (out of 10) for each object under each condition. During each trial, the pneumatic actuation was applied over approximately 2–3 s for the grasp attempt, allowing the actuator to establish contact with the object and apply grasping force. Representative examples from both simulations and experiments are provided in the Supporting Information Video.



**FIGURE 4** | Comparison of experimental and simulated grasping results for the egg (left) and the cylinder (right). The upper tables report the number of successful grasps across different actuation pressure levels and added weight conditions. G1, G2, and G3 denote three independent experimental sets, each consisting of 10 grasping trials conducted under identical conditions. The lower plots show the maximum vertical ( $z$ -axis) position of the simulated object relative to the initial placement ( $z=0$ ), as a function of the applied weight under different actuation pressures.

The success rate comparisons are shown in the upper tables of Figure 4. Overall, the simulation tends to underestimate the actual grasping performance, which is expected as certain inter-modular effects are not captured by the model. Nevertheless, the simulated trends are in good agreement with the experimental results.

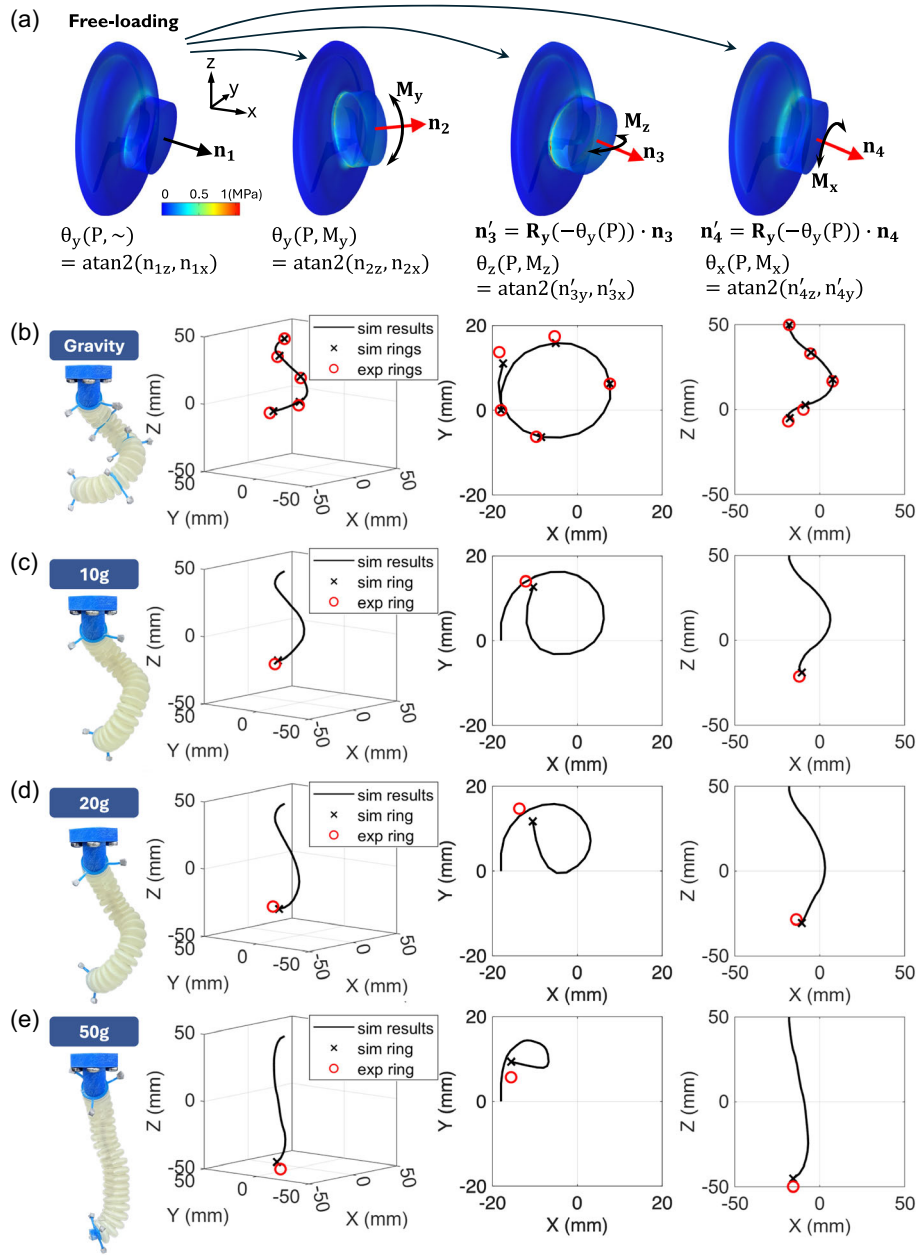
Additionally, the lower plots in Figure 4 show that, although the success rates are underestimated in simulation, the simulated maximum  $z$ -position  $z_{\max}$  of the object correlates with the experimental performance. A higher simulated  $z_{\max}$  is associated with a higher success rate in physical experiments, suggesting it can serve as a useful performance indicator.

#### 4.1.2 | Continuous Deformation With Small NN Surrogates

To evaluate the capability of the surrogate models constructed with small NNs in task-level simulation, and to further validate the proposed pipeline on a representative class of continuous deformation tasks distinct from contact-dominated grasping, we simulated an elephant-trunk-inspired helical soft actuator, adapted from prior work and incorporating a central thread to constrain elongation [25], under gravity and externally applied tip loads.

As shown in Figure 5a, FEM simulations were performed to characterize the deformation of a half-module under two conditions: (i) free-loading pressurization with increasing internal pressure, and (ii) constant pressure combined with externally applied moments about each principal axis. When moments were applied around the nonactive axes ( $z$  and  $x$ ), the rotation about the  $y$ -axis was first reversed to isolate angular deflections  $\theta_z$  and  $\theta_x$ . In this case, the internal pressure is sampled from 0 to 15 kPa with an interval of 0.5 kPa. The externally applied moments are sampled independently along each axis, with  $M_y$  ranging from 0 to  $0.8 \text{ N} \cdot \text{m}$  at intervals of  $0.05 \text{ N} \cdot \text{m}$ , and  $M_x$  and  $M_z$  ranging from 0 to  $0.006 \text{ N} \cdot \text{m}$  at intervals of  $0.0002 \text{ N} \cdot \text{m}$ . Following this axis-wise loading and deformation extraction procedure, three NNs were trained to capture the mappings between pressure, moment, and angular displacement about each axis, each achieving  $R^2 > 0.999$ . This axis-wise modeling choice constitutes a simplification that reduces learning complexity and improves training stability. Cross-axis coupling effects may become more pronounced under larger deformations or combined multi-axis loading.

To validate the surrogate model, we simulated actuator deformation in PyBullet and compared the results with real-world experiments under self-weight (Figure 5b) and additional tip loads of 10, 20, and 50 g (Figure 5c–e). In PyBullet, the actuator was



**FIGURE 5** | Validation of the surrogate model based on small NNs using a helical actuator under gravity and various tip loads. (a) FEM model of the bellow-SPA used in the helical actuator for all three axes. (b) Comparison of the actuator deformation between simulation and averaged physical experiments over three trials under gravity at five ring positions. (c–e) Comparison of the actuator deformation between simulation and averaged physical experiments over three trials under tip loads of 10, 20, and 50 g applied at the actuator tip.

modeled as a chain of links connected by spherical joints, with geometry and mass properties matched to the corresponding SPA modules. The tip load was simulated as a downward external force applied to the terminal link. Details of the ablation study used for selecting the helical actuator simulation parameters, as well as the associated computational cost per simulation run, are provided in the Supporting Information 2.2. Recordings of the simulations are provided in the Supporting Information Video.

Experimental deformation data were recorded using a Motive OptiTrack system (NaturalPoint, Inc., USA) with four Flex 3 cameras, tracking the positions of five rings (each comprising three reflective markers) along the actuator, following the procedure in [25]. The actuation pressure was set to 15 kPa and

supplied using the same pneumatic setup as in Section 4.1.1. This pressure level was empirically identified as the maximum representative operating pressure for this actuator. As the surrogate is trained on FEM data spanning 0–15 kPa, its predictions are expected to remain accurate within this range as an interpolation problem, with validation performed at the upper operating limit.

Each experiment was repeated three times to compute average ring positions. To ensure quasi-static behavior during these measurements, we waited approximately 20–30 s after each actuation step before recording the actuator configuration. For comparison, the actuator base position in PyBullet was aligned with the experimentally measured mean position of the base ring.

Under self-weight, the relative errors at successive ring positions were 1.52%, 1.08%, 2.65%, and 2.85% of the actuator length, with a maximum experimental standard deviation of 1.29 mm observed at the end ring. Under additional tip loads, the relative tip errors were 2.97%, 4.25%, and 5.31% for 10, 20, and 50 g loads, respectively, with a maximum experimental standard deviation of 1.33 mm at 20 g.

### 4.1.3 | A Tendon-Driven Finger Simulation

To demonstrate the applicability of the proposed modeling approach beyond bellow actuators and pneumatic actuation, we also applied the NN surrogate model to a tendon-driven actuator from prior work [28], simulating its deformation under free loading and tip loads of 20 and 50 g with the same tendon travel.

As shown in Figure 6a, FEM simulations were conducted to characterize the deformation  $\theta$  of an actuator module (as indicated by the dashed box in Figure 6e) under free loading and externally applied moments  $M_{ext}$ , with the tendon force modeled as a directional load following the tendon path. In this case, the tendon force is sampled from 0 to 8 N with an interval of 0.05 N. The externally applied moment is sampled from 0 to 0.05 N · m at intervals of 0.0025 N · m. A NN was trained to capture the mapping between tendon force, external moment, and angular deflection, achieving  $R^2 > 0.999$ . In addition, a simple regression was used to map tendon force to tendon travel.

To validate the sim-to-real transfer, the tendon-driven actuator was modeled in PyBullet as a chain of links connected by revolute joints, with geometry and mass properties taken from the CAD model. The tip load was simulated as an external force applied to the terminal link and directed toward the pulley position used in the real experiments, which corresponds to the finger's initial tip position. Simulation results were obtained by fine-tuning the tendon force to reproduce the tendon travel reported in [28]. The PyBullet simulation parameters for this tendon-driven actuator were chosen to be identical to those used in

the helical actuator experiments, as no object contact was involved. Recordings of the simulations are provided in the Supporting Information Video.

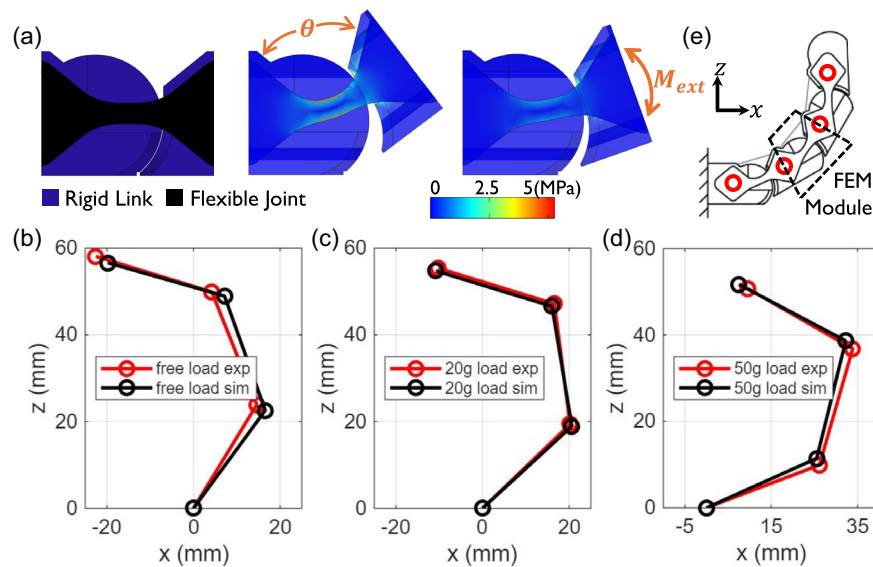
The simulation results show good agreement with experimental measurements reported in [28], where each loading condition was experimentally repeated three times, with a maximum standard deviation of 3.51 mm at the end link position under free load. As illustrated in Figure 6b–d, the resulting tip errors were 3.14 mm under free loading, 0.84 mm under a 20 g tip load, and 2.29 mm under a 50 g tip load.

## 4.2 | Surrogate Models for Task-Driven Designs

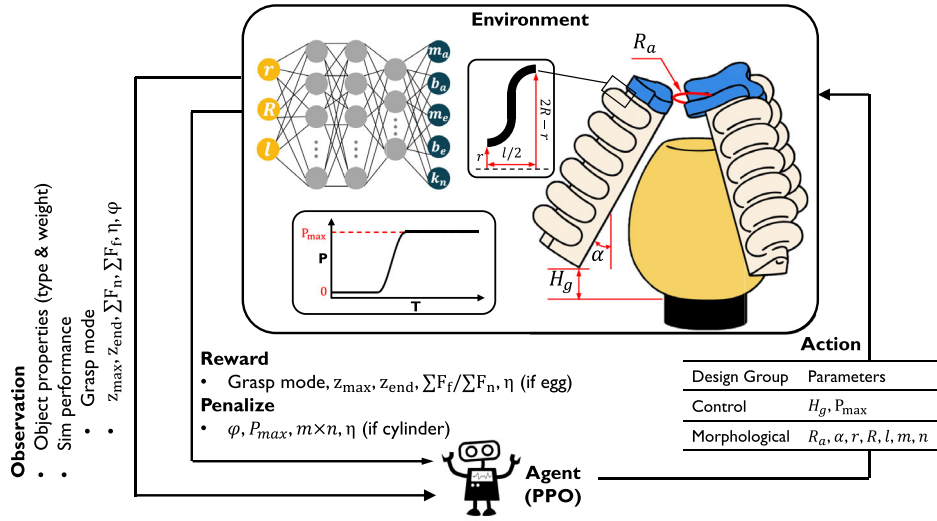
In this section, we evaluated the effectiveness of the proposed modeling pipeline for task-driven design by coupling the meta-models with SOTA optimization algorithms. Two representative design studies are considered: (i) soft gripper co-design using reinforcement learning (RL), and (ii) 3D actuator shape matching using evolutionary optimization. These two cases employ different meta-model instantiations and optimization strategies, thereby demonstrating that the proposed pipeline can flexibly incorporate task-specific surrogate representations and remain fully compatible with modern learning-based and gradient-free optimization methods.

### 4.2.1 | Designing Soft Gripper With RL

The first case study employs the polynomial-based meta-model in an RL-driven soft gripper co-design setting, as illustrated in Figure 7. Despite its simplified structure, the polynomial surrogate has been shown to provide sufficient accuracy and smoothness for the considered bellow-SPA-based soft gripper within the investigated operating range (Section 4.1.1), while maintaining the computational efficiency required for RL-based design exploration. Here, both morphological design parameters and actuation-related variables are optimized jointly through task-level



**FIGURE 6** | Validation of the surrogate modeling approach beyond the bellow-SPA. (a) FEM model of the joint in the tendon-driven finger developed by [28]. (b–d) Comparison of link positions between simulation and averaged experimental results under free load, 20, and 50 g tip loads, respectively. Circle markers correspond to the center positions of each rigid link, with the base link center used as the origin of the reference frame, as indicated in (e), adapted with permission from [28]. Copyright 2025 IEEE.



**FIGURE 7** | Schematic overview of task-driven soft gripper design using RL. An RL agent (PPO) is trained on a task distribution consisting of grasping an egg and a cylinder with varying weights. The environment corresponds to a simulated grasping trial with randomized object type and weight, where both control and morphological parameters are sampled from the action space and evaluated using the meta-model for polynomial surrogates. The agent observes task properties and simulated grasping performance features, and assigns rewards and penalties accordingly to learn a policy that maps object-specific conditions to optimal gripper design parameters.

grasping performance, with the meta-model providing fast predictions of module-level mechanical responses during simulation.

We first constructed the meta-model from a family of polynomial-based surrogates. The same polynomial surrogate modeling procedure as described in Section 4.1.1 was applied repeatedly over a discretized design space defined by the actuator inner radius  $r \in \{2.75, 3.5, 4.25, 5\}$  mm, wall thickness  $t = 1.5$  mm, average radius  $R \in \{6.75, 7.5, 8.25, 9\}$  mm, and module length  $l \in \{9, 10, 11, 12\}$  mm. This design space is chosen as a local neighborhood around the baseline gripper design used in Section 4.1.1, enabling task-driven optimization within a practically relevant actuator configuration range while avoiding unrealistic actuator geometries that would violate the manufacturability and geometric constraints of the bellow-SPA design. This resulted in 54 valid actuator designs, with the remaining combinations excluded by the geometric boundary conditions imposed on the design parameters [26].

A meta-model was then trained as described in Section 3.4, using a three-dimensional input vector composed of the design parameters  $(r, R, l)$ . The network comprises three hidden layers, two with 64 neurons and one with 32 neurons, each using ReLU activation, followed by a linear output layer predicting the five polynomial coefficients  $(m_a, b_a, m_e, b_e, k_n)$ , as defined in the polynomial surrogate formulation detailed in the Supporting Information 1.3, Equation (S7). The model was trained with the Adam optimizer to minimize the mean squared error (MSE) for up to 1000 iterations and achieves an  $R^2$ -score exceeding 0.99 on the full dataset.

This meta-model was embedded into the PyBullet simulation to replace explicit per-design polynomial surrogates, enabling the prediction of the actuator’s grasping response for continuously varying design parameters. The actuator geometries corresponding to different design parameters are generated by scaling the baseline STL mesh used in Section 4.1.1. The task-level

simulation takes three groups of inputs. The control domain contains the grasping height  $H_g$  (distance between the actuator tip and the object bottom) and the maximum pressure  $P_{\max}$ . The morphological domain includes the actuator base radius  $R_a$  (radius of the circle formed by the actuator bases), inclination angle  $\alpha$  (angle between the actuator longitudinal axis and gravity), the geometric parameters  $(r, R, l)$ , the module number  $n$  (modules per actuator) and the actuator number  $m$  (number of actuators). The object properties comprise the object type (“egg” or “cylinder”) and the weight number (the number of attached 20 g weights). All other simulation settings follow those in Section 4.1.1.

The simulator outputs a set of task-relevant features: the maximum and final object positions along the vertical axis  $z_{\max}, z_{\text{end}}$ ; the summed normal and friction contact forces over the grasping stage  $\sum F_n, \sum F_f$ ; the contact module efficiency  $\eta$  (ratio between the number of active contact modules at the end of grasping and the total module count); the slip ratio  $\varphi$  (ratio between the average object velocity and gripper velocity along the  $z$ -axis); and the grasping mode (miss, touch, slip, or stable grasp) determined from object motion and contact forces.

On top of this simulation, we constructed an RL environment for task-driven gripper design. The action space contains all control and morphological design variables listed above, while the observation space includes the object properties together with the simulated grasping performance features. The reward function encourages successful grasping modes, higher lifting height, and larger friction-to-normal force ratios, while penalizing excessive pressure and material usage. It further promotes distributed multi-module contacts when grasping the egg and more localized contacts when grasping the cylinder.

Task-driven design was then performed using a Proximal Policy Optimization (PPO) agent trained on a task distribution consisting of grasping an egg and a cylinder with varying weights. Each episode corresponds to a single simulated grasp attempt, with the

object type and weight randomized at reset, so that the learned policy maps task-dependent object conditions to both control commands and gripper morphology. Training was run for 14 400 total timesteps in a vectorized environment with four parallel simulations, using the hyperparameters detailed in Supporting Information 2.3. The associated learning curves, shown in Figure S8, indicate stable convergence of the policy and value function.

The resulting distributions of the optimized parameters are summarized in Figures S6 and S7, obtained by evaluating the trained policy on each object type–weight combination and sampling the predicted action parameters over 10 stochastic rollouts per task condition. Most morphological design variables, including inner radius, average radius, module length, module number, and actuator number, converge to nearly identical values across all tasks, whereas the actuator base radius exhibits a clear task-dependent differentiation between the egg and cylinder. This suggests that, under the current reward and task distribution, the policy predominantly converges to a robust general-purpose gripper morphology, with a small subset of parameters adapting to object geometry.

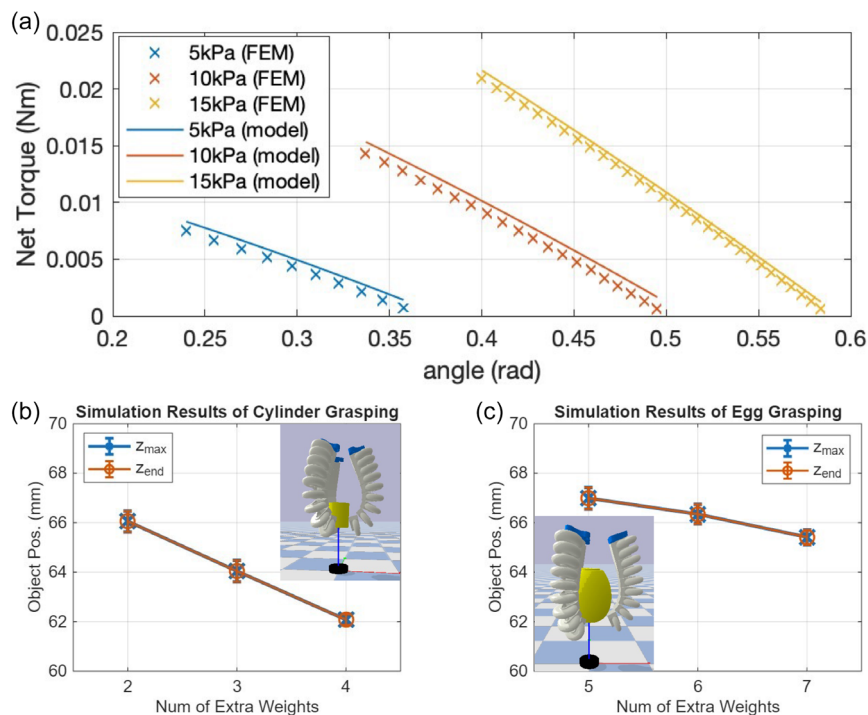
To further verify the quality of the learned design and to remove potential uncertainty introduced by the meta-model, we recomputed the polynomial surrogate for the final optimized design parameters using direct FEM data. Notably, this design does not belong to the original meta-model training set. The reconstructed surrogate achieved a root-mean-squared error (RMSE) of 3.21% relative to the maximum net torque value when compared with FEM data (Figure 8a). Using this refined surrogate, grasping simulations (10 attempts with added environmental

noise) were carried out for all task conditions. The simulation protocol follows the same procedure used in Section 4.1.1 for the sim-to-real validation of the baseline gripper. As shown in Figure 8b and c, the simulated maximum z-position  $z_{\max}$  and final z-position  $z_{\text{end}}$  of the object are nearly identical when grasped by the optimized gripper, indicating stable grasps and corresponding to a success rate of 10/10 in all trials. A quantitative comparison with the baseline gripper configuration reported in Section 4.1.1 is provided in the Supporting Information Table S1, which summarizes the grasp success rates and the simulated maximum object position  $z_{\max}$  across the tested task conditions. These results confirm that the proposed pipeline can exploit RL together with the polynomial-based meta-model to support task-driven soft gripper design optimization.

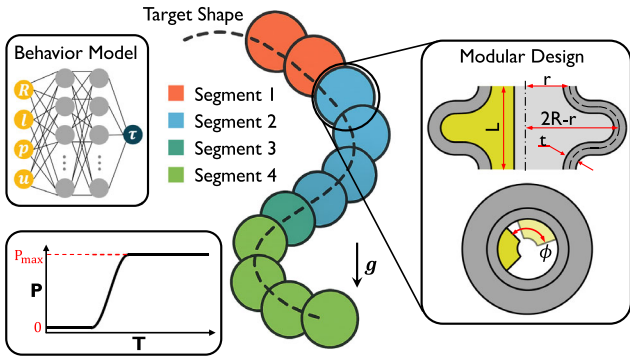
#### 4.2.2 | 3D Shape-Matching With CMA-ES

The second case study uses the NN-based meta-model within a representative evolutionary optimization method for actuator shape-matching design, as shown in Figure 9. The meta-model predicts the deformation behavior of candidate designs, while a gradient-free optimizer updates the design parameters according to a geometric matching objective. In this case study, the shape-matching task is defined using the same right-handed helical target shape as in [25], with a radius of 18 mm and a height of 60 mm.

We first applied the NN surrogate model to the same actuator design for deformation prediction along the three principal axes as described in Section 4.1.2. This procedure was repeated over a discretized design space defined by the average radius  $R \in \{5, 5.5, 6, 6.5, 7\}$  mm and module length  $l \in \{4, 4.5, 5, 5.5, 6\}$



**FIGURE 8** | Results of the RL-optimized design. (a) Comparison of net torque versus angular deflection under representative pressure levels between the polynomial surrogate model for the optimized design and its FEM simulation, showing close agreement consistent with the validation in Section 4.1.1. (b) Simulated object grasping results using the polynomial surrogate model for the optimized design, indicating successful and stable grasping under the intended task conditions.



**FIGURE 9** | Schematic overview of the shape-matching actuator design task. The actuator is composed of multiple segments (highlighted in different colors), each characterized by the number of identical modules, the geometric design parameters of the module, and the rotation angle relative to the preceding segment. Behavior models that map the modular design and actuation inputs to the output response are employed in simulation to obtain the deformed shape of the designed actuator under actuation pressure and gravity, in order to match a given target shape.

mm, while keeping the inner radius ( $r = 3$  mm) and wall thickness ( $t = 1.5$  mm) fixed. This design space is centered around the optimal actuator configuration reported in [25], allowing interpolation-based evaluation of the meta-model in the vicinity of a known high-performance design. After excluding combinations that violate the geometric condition  $R - r \geq l/4$  [25], 24 valid actuator designs remained.

Following Section 3.4, a behavior model was trained for each axis by aggregating data from all sampled designs. For a design characterized by  $(R, l)$ , all tuples  $(p, u, \tau)$ , corresponding to actuation pressure, half joint angle, and resulting joint torque, were collected and stacked into input-output pairs of the form  $(R, l, p, u) \rightarrow \tau$ . Each axis-wise behavior model was implemented as a feedforward NN with two hidden layers of 64 neurons and ReLU activation, followed by a linear output layer predicting the net joint torque. The models were trained with the Adam optimizer to minimize MSE, achieving  $R^2$  scores exceeding 0.999 for all three axes.

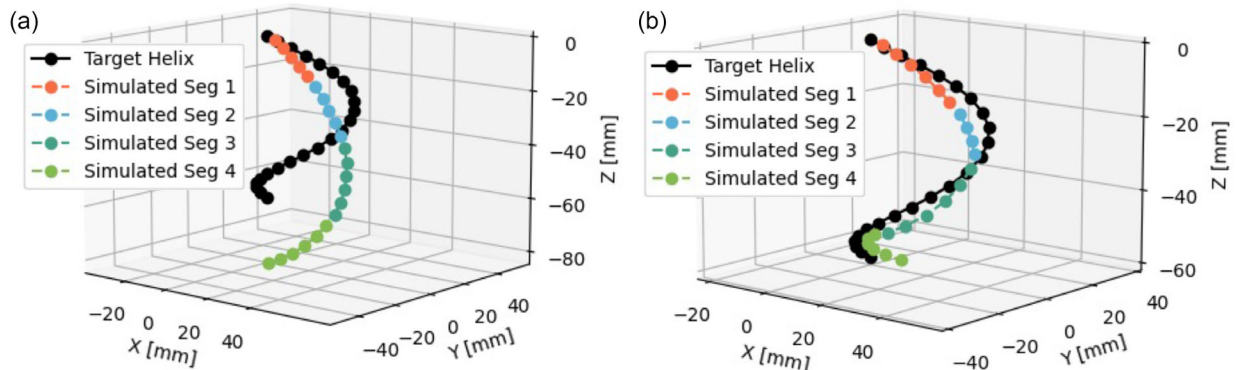
These design-conditioned behavior models were then embedded into the PyBullet simulation environment to replace the explicit

small NN models, enabling fast prediction of actuator deformation for varying design parameters. The simulation takes two groups of inputs. The first defines the actuator design: the number of segments  $N$ ; the average radius  $R$  and module length  $l$  shared by all modules within each segment; the number of modules per segment  $n$ ; the relative rotation angle between adjacent segments  $\phi$ , and the actuation pressure  $P_{\max}$ . The second group specifies the target shape to be matched.

For each candidate design and pressure, the PyBullet simulation produces the deformed centerline of the actuator, reconstructed as a piecewise arc-length-parameterized curve based on link positions and joint rotations. The target helix is resampled according to the normalized cumulative arc length of this curve to ensure one-to-one correspondence between simulated and target points. Shape-matching quality is quantified using the RMSE and the maximum point-wise Euclidean distance  $e_{\max}$  between corresponding points, and these metrics serve as objective functions for design optimization. In [25], no satisfactory actuator could be identified due to limitations of the underlying model, with the best design yielding an RMSE of 25.68 mm and an  $e_{\max}$  of 36.20 mm (Figure 10a).

The optimization problem is formulated as a joint search over structural parameters, geometric design parameters, and actuation pressure. At the structural level, the actuator is partitioned into  $N = 4$  segments in this study. The total number of modules is first bounded using a physically motivated estimate based on the target curve length and the range of module lengths, and then five candidate totals are sampled around the center of this interval. For each total, multiple structural configurations  $n_{ii=1}^N$  are generated by distributing modules across segments under a balance constraint, avoiding extreme structural imbalance while preserving diversity. This procedure yields 81 structural candidates.

For each structural candidate, the continuous design parameters  $(R_i, l_i, \phi_i)_{i=1}^N$  and actuation pressure  $p$  are optimized jointly. To remove global rotational redundancy, the first intersegment rotation is fixed to  $\phi_1 = 0$ , while the remaining  $\phi_i$  are treated as optimization variables. Continuous optimization is carried out using the Covariance Matrix Adaptation Evolution Strategy (CMA-ES), which is well suited to strongly nonlinear, nonconvex objective landscapes and black-box settings where each function evaluation involves a full physics-based simulation with embedded



**FIGURE 10** | Results of the shape-matching design optimization. (a) Comparison between the target helix and the optimized actuator obtained using the previous modeling approach [25]. (b) Comparison between the target helix and the optimized actuator obtained using the proposed behavior models.

surrogate models. The initial step size is set to  $\sigma_0 = 0.3$  to balance global exploration and local refinement.

During the shape-matching process, CMA-ES iteratively searches the continuous design space for each of the 81 structural configurations, minimizing both RMSE and  $e_{\max}$  until either convergence below 5 mm for both metrics or 60 iterations are reached. The resulting optimal actuator design and pressure are reported in Supporting Information 2.4, and summarized in Table S2. Based on the behavior model prediction, the optimized design achieved an RMSE of 4.51 mm and an  $e_{\max}$  of 8.46 mm with respect to the target helix.

To further evaluate the fidelity of the behavior model prediction, the small NN surrogate models for each segment were subsequently reconstructed using direct FEM data. All reconstructed surrogates achieved  $R^2$  scores above 0.999. The corresponding actuator was then simulated again, yielding an RMSE of 5.05 mm and  $e_{\max} = 8.19$  mm with respect to the target helix (Figure 10b). Compared with the previous best result, this represents a substantial improvement in shape matching and confirms that the NN-based behavior models, when combined with a representative evolutionary optimizer such as CMA-ES, are effective for task-driven design within the proposed pipeline.

### 4.3 | General Discussion and Limitations

Building on the results presented above, this subsection provides a brief general discussion of the limitations underlying the proposed approach.

First, the proposed pipeline primarily supports interpolation within parameterized actuator design spaces that are explicitly sampled during surrogate construction, and does not aim to generalize across fundamentally different morphologies or actuation principles. Second, the surrogate models were constructed from quasi-static FEM data, which neglect dynamic effects such as rate-dependent material behavior and transient interactions. Third, while extensive validation was performed mainly within the sampled ranges, systematic evaluation outside these ranges remains limited. Finally, hardware validation was demonstrated for selected representative cases, but not for all actuator types and task scenarios considered in simulation.

## 5 | Conclusion

This paper introduced a unified FEM-based surrogate modeling pipeline for generalized task-driven design of soft robots. By integrating high-fidelity FEM data generation, compact surrogate modeling, design-conditioned meta-model learning, and PRBM-based task-level simulation, the proposed pipeline bridges the long-standing gap between physical accuracy and computational efficiency in soft robotic design. Through extensive sim-to-real validation across multiple actuator types and two representative task-driven design studies, namely RL-based soft gripper co-design and evolutionary shape-matching optimization, we demonstrated that the pipeline supports both reliable physical prediction and efficient task-level optimization under realistic environmental interactions. Beyond the specific tasks investigated in this work, the modular structure of the proposed pipeline facilitates extension to additional soft robotic applications

within similar modeling assumptions. In particular, the reliance on physically grounded FEM data and surrogate abstraction makes the pipeline well-suited for environments involving complex interactions, such as underwater manipulation, fluid-structure coupling, and confined-space operations. Moreover, the decoupling between surrogate modeling and task-level optimization allows the proposed approach to be readily combined with diverse learning-based and gradient-free optimization paradigms. These properties position the pipeline as a fundamental building block toward scalable, autonomous, and general-purpose task-driven design of soft robotic systems in complex real-world environments.

---

### Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Grant (EP/V000748/1).

### Funding

This study was supported by Engineering and Physical Sciences Research Council (EP/V000748/1).

### Conflicts of Interest

The authors declare no conflicts of interest.

### Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

### References

1. D. Rus and M. T. Tolley, "Design, Fabrication and Control of Soft Robots," *Nature* 521 (2015): 7553–7467.
2. C. Lee, M. Kim, Y. J. Kim, et al., "Soft Robot Review," *International Journal of Control, Automation and Systems* 15, no. 1 (2017): 3.
3. J. Shintake, V. Cacucciolo, D. Floreano, and H. Shea, "Soft Robotic Grippers," *Advanced Materials* 30, no. 29 (2018): 1707035.
4. J. Pinskiel and D. Howard, "From Bioinspiration to Computer Generation: Developments in Autonomous Soft Robot Design," *Advanced Intelligent Systems* 4, no. 1 (2022): 2100086.
5. F. Stella and J. Hughes, "The Science of Soft robot Design: A Review of Motivations, Methods and Enabling Technologies," *Frontiers in Robotics and AI* 9 (2023): 2022, <https://doi.org/10.3389/frobt.2022.1059026>.
6. M. Pozzi, E. Miguel, R. Deimel, et al., *2018 IEEE International Conference on Robotics and Automation (ICRA)*, (2018), 4206–4213.
7. M. S. Xavier, A. J. Fleming, and Y. K. Yong, "Finite Element Modeling of Soft Fluidic Actuators: Overview and Recent Developments," *Advanced Intelligent Systems* 3, no. 2 (2021): 2000187.
8. C. Schaff, A. Sedal, S. Ni, and M. R. Walter, "Sim-to-Real Transfer of Co-Optimized Soft Robot Crawlers," *Autonomous Robots* 47, no. 8 (2023): 1195–1211.
9. T. Navez, E. Ménager, P. Chaillou, O. Goury, A. Kruszewski, and C. Duriez, "Modeling, Embedded Control, and Design of Soft Robots Using a Learned Condensed FEM Model," *IEEE Transactions on Robotics* 41 (2025): 2441.
10. M. S. Xavier, C. D. Tawk, A. Zolfagharian, et al., "Soft Pneumatic Actuators: A Review of Design, Fabrication, Modeling, Sensing, Control and Applications," *IEEE Access* 10 (2022): 59442.

11. C. Alessi, C. Agabiti, D. Caradonna, C. Laschi, F. Renda, and E. Falotico, "Rod Models in Continuum and Soft Robot Control: A Review," 2024, <https://arxiv.org/abs/2407.05886>.
12. D. Bruder, M. A. Graule, C. B. Teeple, and R. J. Wood, "Increasing the Payload Capacity of Soft Robot Arms by Localized Stiffening," *Science Robotics* 8, no. 81 (2023): eadf9001.
13. R. J. Webster III and B. A. Jones, "Design and Kinematic Modeling of Constant Curvature Continuum Robots: A Review," *The International Journal of Robotics Research* 29, no. 13 (2010): 1661.
14. H. Zhang, K. H. Heung, G. J. Naquila, A. Hingwe, and A. D. Deshpande, "A Novel Parameter Estimation Method for Pneumatic Soft Hand Control Applying Logarithmic Decrement for Pseudo-Rigid Body Modeling," *Advanced Intelligent Systems* 7, no. 3 (2025): 2400637.
15. C. Della Santina, R. K. Katzschmann, A. Biechi, and D. Rus, 2018 *IEEE International Conference on Soft Robotics (RoboSoft)*, (2018), 46–53.
16. M. A. Graule, C. B. Teeple, T. P. McCarthy, G. R. Kim, R. C. St. Louis, and R. J. Wood, 2021 *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, (2021), 3934–3941.
17. P. Moseley, J. M. Florez, H. A. Sonar, G. Agarwal, W. Curtin, and J. Paik, "Modeling, Design, and Development of Soft Pneumatic Actuators with Finite Element Method," *Advanced Engineering Materials* 18, no. 6 (2016): 978.
18. P. Ferrentino, E. Roels, J. Brancart, S. Terryn, G. Van Assche, and B. Vanderborght, "Finite Element Analysis-Based Soft Robotic Modeling: Simulating a Soft Actuator in SOFA," *IEEE Robotics & Automation Magazine* 31, no. 3 (2024): 97.
19. M. Macklin, M. Müller, and N. Chentanez, *Proceedings of the 9th International Conference on Motion in Games, MIG '16* (Association for Computing Machinery, 2016), 49–54, <https://doi.org/10.1145/2994258.2994272>.
20. Y. Hu, Y. Fang, Z. Ge, et al., "A Moving Least Squares Material Point Method with Displacement Discontinuity and Two-Way Rigid Body Coupling," *ACM Transactions on Graphics* 37 (2018): 4.
21. M. Wiese, B.-H. Cao, and A. Raatz, 2022 *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, (2022), 9371–9378.
22. O. Goury and C. Duriez, "Fast, Generic, and Reliable Control and Simulation of Soft Robots Using Model Order Reduction," *IEEE Transactions on Robotics* 34, no. 6 (2018): 1565.
23. A. T. Mathew, I. B. Hmida, C. Armanini, F. Boyer, and F. Renda, "SoRoSim: A MATLAB Toolbox for Hybrid Rigid-Soft Robots Based on the Geometric Variable-Strain Approach," *IEEE Robotics & Automation Magazine* 30, no. 3 (2023): 106.
24. J. Liu, J. H. Low, Q. Q. Han, et al., "Simulation Data Driven Design Optimization for Reconfigurable Soft Gripper System," *IEEE Robotics and Automation Letters* 7, no. 2 (2022): 5803.
25. Y. Yao, L. He, and P. Maiolino, "SPADA: A Toolbox of Designing Soft Pneumatic Actuators for Shape Matching Based on Surrogate Modeling," *Robotics Reports* 2, no. 1 (2024): 1.
26. Y. Yao, L. He, and P. Maiolino, 2022 *IEEE. 5th International Conference on Soft Robotics (RoboSoft)*, (2022), 29–34.
27. A. A. Shabana, M. A. Abdelrahman, A. M. El-Kassas, and H. E. Abdel-Azim, "Vulval Syringoma: First Report of a Case in an Egyptian Female," *Journal of the Egyptian Society of Tribology* 9, no. 4 (2012): 28, article 3.
28. E. Capp, M. Pontin, P. Walters, and P. Maiolino, 2025 *IEEE. 8th International Conference on Soft Robotics (RoboSoft)*, (2025), 1–8.

## Supporting Information

Additional supporting information can be found online in the Supporting Information section. **Supporting Fig. 1:** FEM setup and force evolution

for a translational bellow-SPA module. (a) The FEM model under free-loading and constrained conditions. (b) Calculated elastic, actuation, and external forces over the simulation steps, illustrating the transition from the free-loading phase to the constrained phase. **Supporting Fig. 2:** Polynomial-based surrogate modeling process and validation. (a–d) Polynomial fits of the actuation and elastic stiffness across different pressures and displacements. (e) FEM model of connected actuator modules. (f) Comparison between predictions from the polynomial surrogate model and the full FEM simulation for actuators composed of 2–10 modules. **Supporting Fig. 3:** Surrogate modeling using small NNs. (a) FEM model of a bending bellow-SPA with an internal constraint. (b) Calculated actuation stiffness as a function of angular deflection under multiple pressure levels, obtained using the formulation in Section 3.2. (c) Architecture of the small feedforward NN. (d) Training and prediction results of the neural network surrogate. **Supporting Fig. 4:** Ablation study of the PyBullet simulation settings for the grasping task. The effects of the simulation time step size and the joint damping coefficient are evaluated under fixed actuation pressure, contact friction, and object properties. Representative actuator joint angle and joint torque responses, together with the maximum vertical position of the object  $z_{\max}$ , are analyzed. **Supporting Fig. 5:** Ablation study of the PyBullet simulation settings for the helical actuator. The effects of the simulation time step size and the joint damping coefficient are evaluated at the fifth ring (actuator tip). (a) Position error between the experimentally measured actuator tip and the simulated tip position, where the simulated position is computed as the mean over the last half cycle after the motion has settled. (b) Standard deviation of the simulated tip position over the last half cycle. **Supporting Fig. 6:** Distributions of the RL-generated design and control parameters for different object types. Histograms show the distributions of the optimized parameters for grasping an egg (blue) and a cylinder (orange), including grasping height, maximum pressure, actuator base radius, inclination angle, inner radius, average radius, module length, module number, and actuator number. The  $y$ -axis represents the number of evaluation samples, while the  $x$ -axis corresponds to the parameter indicated in each subplot title. When two bars appear side by side, they share the same parameter value, defined by the center of the two bars; when only a single bar is present, its center indicates the corresponding parameter value. **Supporting Fig. 7:** Distributions of the RL-generated design and control parameters across object types and weight conditions. Histograms show the optimized parameter distributions for grasping eggs and cylinders under different weight cases, including grasping height, maximum pressure, actuator base radius, inclination angle, inner radius, average radius, module length, module number, and actuator number. The  $y$ -axis represents the number of evaluation samples, while the  $x$ -axis corresponds to the parameter indicated in each subplot title. When bars appear side by side, they share the same parameter value, defined by the center of the bars; when only a single bar is present, its center indicates the corresponding parameter value. **Supporting Fig. 8:** Training performance of the RL agent (PPO). (a) Mean episode reward, (b) value loss, (c) total training loss, (d) explained variance, (e) policy entropy, and (f) approximate Kullback–Leibler (KL) divergence, plotted over the training steps. **Supporting Table 1:** Quantitative comparison between the baseline gripper and the RL-optimized gripper, evaluated using the grasp success rate and the averaged simulated maximum object lifting position  $z_{\max}$ . **Supporting Table 2:** Optimized actuator design parameters for the 3D helix shape-matching task.