

Sim2Real for Soft Robotic Fish via Differentiable Simulation

John Z. Zhang^{1,2}, Yu Zhang¹, Pingchuan Ma³, Elvis Nava^{1,4}, Tao Du³, Philip Arm¹,
Wojciech Matusik³, Robert K. Katzschmann^{1,4}

Abstract—Accurate simulation of soft mechanisms under dynamic actuation is critical for the design of soft robots. We address this gap with our differentiable simulation tool by learning the material parameters of our soft robotic fish. On the example of a soft robotic fish, we demonstrate an experimentally-verified, fast optimization pipeline for learning the material parameters from quasi-static data via differentiable simulation and apply it to the prediction of dynamic performance. Our method identifies physically plausible Young’s moduli for various soft silicone elastomers and stiff acetal copolymers used in creation of our three different robotic fish tail designs. We show that our method is compatible with varying internal geometry of the actuators, such as the number of hollow cavities. Our framework allows high fidelity prediction of dynamic behavior for composite bi-morph bending structures in real hardware to millimeter-accuracy and within 3% error normalized to actuator length. We provide a differentiable and robust estimate of the thrust force using a neural network thrust predictor; this estimate allows for accurate modeling of our experimental setup measuring bollard pull. This work presents a prototypical hardware and simulation problem solved using our differentiable framework; the framework can be applied to higher dimensional parameter inference, learning control policies, and computational design due to its differentiable character.

I. INTRODUCTION

Soft robots are favored over traditional hard robots in a growing list of scenarios, including bio-inspired design, conformal gripping, co-bot situations, and more. While the design and simulation of hard robots are well understood and extensively shown thanks to decades of development, the same simulation task for compliant structures and their interaction with the environment is an ongoing challenge.

Although many simulation schemes exist, none have demonstrated sufficient sim2real matching through verifiable experiments for soft robots in a hydrodynamic situation such as a robotic fish. The high dimensionality of soft mechanisms and the complex physics of fluid-structure interactions (FSI) make this problem particularly difficult. One critical challenge in producing high accuracy simulations is the precise measurement and manual tuning needed to dial-in the material parameters of soft mechanisms. The problem has recently become more tractable because of advances in physics-based simulation, such as differentiability and data-driven approaches.

*We are grateful for funding received by the ETH AI Center, and the Defense Advanced Research Projects Agency.

¹Soft Robotics Lab, ETH Zurich, Switzerland

²Mechanical Engineering, MIT, USA

³Computer Science and AI Lab, MIT, USA

⁴ETH AI Center, ETH Zurich, Switzerland

Corresponding author: Robert K. Katzschmann, rkk@ethz.ch

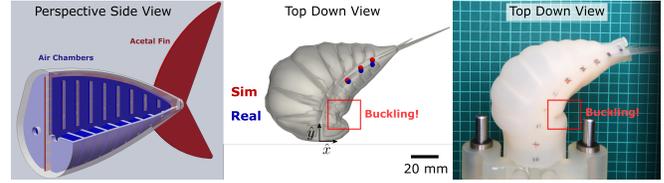


Fig. 1. *Left*: A model of a pneumatic fish tail (*Nemo*). *Middle*: The same fish tail inflated to 250 mbar in simulation. *Right*: The actual hardware of the fish tail in an experimental setup for measuring the deformation and the thrust of the fish in air and underwater. Our framework is capable of reproducing the buckling effect in simulation, a valuable result.

We present a differentiable simulation framework that can be used to accurately predict the deformation of a soft pneumatic actuator and identify its material properties through gradient-based optimization. Three designs of a pneumatically-actuated fish tail are used as benchmark (see Table I) since composite bi-morph structures are difficult to simulate due to the disparate Young’s moduli and aspect ratios of the deformable body and stiffer spine. The deformable body of the fish tail is cast out of silicone elastomer and has a flexible spine made of an acetal plastic sheet in the center (see Fig. 1).

Our framework is capable of learning physically-plausible material parameters, i.e., Young’s modulus and Poisson’s ratio, using only a quasistatic data set without extensive material testing. We demonstrate that after optimization we can produce accurate simulation results compared to dynamic data collected in our hardware setup to within millimeter accuracy or within 3% max error normalized to the actuator length. The fast system identification is achieved using our finite element method (FEM) [1], [2], which combines projective dynamics [3] and differentiable simulation [4].

To verify our simulation results, we developed an experimental setup for collecting position, pressure, and force data using marker tracking, a pneumatic valve array, and a load cell. The raw data is synchronized and analyzed in MATLAB.

Furthermore, we are able to reproduce the measured thrust in our bollard-pull type experiment using a simple neural network predictor that can be integrated in our differentiable simulator. Our model is capable of predicting thrust performance in actuation signals not seen in the training set and preserves differentiability of the simulator. Our aim is to provide straightforward sim2real methods for roboticists that are within reach: such methods could be used confidently for design optimization if they are verifiable by real data.

In this paper, we contribute:

- a simulation framework for a soft pneumatic bending actuator consisting of disparate materials and geometries;
- a system identification method that uses differentiable simulation and gradient-based optimization to accurately learn material parameters of two isotropic corotated materials;
- a data-driven hydrodynamics model using a neural network as a simple predictor of the thrust force generated by the fish; and
- experimental verification of our simulation results using a hardware setup suitable for data acquisition in both air and water.

II. RELATED WORK

A. *Soft underwater robots*

Soft robots are difficult to optimally design and control when compared to their rigid counterparts due to the infinite dimensionality of their compliant structures. Due to this modeling complexity, an experienced designer must hand craft each design guided by intuition, experiments, and approximate models. Marchese et al. offer approaches to designing and fabricating soft fluidic elastomer robots, the type of robot we are also using in this work [5]. Katzschmann et al. present the design, fabrication, control, and testing of a soft robotic fish with interior cavities that is hydraulically actuated. Their manually designed robot can swim at multiple depths and record aquatic life in the ocean [6], [7]. Zhu et al. manually optimize the swimming performance of their robotic fish, Tunabot [8]. The authors measured kinematics, speed, and power at increasing flapping frequencies to quantify swimming performance and find agreement in performance between real fish and their Tunabot over a wide range of frequencies. Zheng et al. propose to design soft robots by pre-checking controllability during the numerical design phase [9]. FEM is used to model the dynamics of cable-driven parallel soft robot and a differential geometric method is applied to analyze the controllability of the points of interest. Katzschmann et al. [10] manually tweak the material parameters of their reduced-order FEM [11] with an experimental soft robotic arm to perform dynamic closed-loop control. Van et al. present a DC motor driven soft robotic fish which is optimized for speed and efficiency based on experimental, numerical and theoretical investigation into oscillating propulsion [12]. Wolf et al. use a pneumatically-actuated fish-like stationary model to investigate how parameters like stiffness, strength, and frequency affect thrust force generation [13]. Wolf et al. measure thrust, side forces, and torques generated during propulsion and use a statistical linear model to examine the effects of parameter combinations on thrust generation; they show that both stiffness and frequency substantially affect swimming kinematics. We are not aware of any work that uses a fast differentiable FEM simulation environment to learn material parameters for soft robotic fish using a bollard-pull style experimental setup.

B. *Differentiable soft-body simulators*

Our work is also relevant to the recent developments of robotic simulators, particularly for soft robots. Geilinger et al. [14] present a differentiable multi-body dynamics solver that is able to handle frictional contact for rigid and deformable objects. Coevoet et al. [15] notably present a non-differentiable framework for modeling, simulation, and control of soft-bodied robots using continuum mechanics for modeling the robotic components and using Lagrange multipliers for boundary conditions like actuators and contacts. Most related to our work are the recent works on differentiable soft-body and fluid simulators [16], [1], [17], [4], [18], [19], [2]. These papers develop numerical methods for computing gradients in a traditional simulators. Furthermore, they demonstrate the power of gradient information in robotics applications, e.g., system identification or trajectory optimization. Most of the works present simulation results only, with ChainQueen [18] and Real2Sim [17] being two notable exceptions that discuss real-world soft-robot applications. Notably, [17] optimizes visco-elastic material parameters of a finite element simulation to approximate the dynamic deformations of real-world soft objects, such as an open-loop controlled tendon-driven crawling robot. Bern et al. [20] have also demonstrated the use of differentiable simulation to learn from a quasi-static data set for the purpose of optimizing open-loop control inputs. Dubied et al. [21] is the most recent example that demonstrates sim2real matching for a soft robotic fish tail, shows system identification on a passive structure for just the Young's modulus, and investigates the mismatch in damping between reality and simulation. In this previous work, the fish tail actuation is simulated using a simplified muscle model and only one design is shown whereas in this paper, the pressure boundary condition is simulated exactly as fabricated for each pneumatic chamber geometry for three different designs. Simulating the pneumatic chambers improves accuracy and allows for physically-plausible Young's moduli and Poisson ratios to be identified. In this current work, we further demonstrate that the gradient-based optimization can be carried out to higher dimensional design spaces that include more than one material parameter.

C. *Hydrodynamic Surrogates*

For underwater soft robots, the challenge of simulation is exacerbated by the hydrodynamic interaction with the soft body. Several previous works tackle the fluid-structure interaction problem through different methods, including heuristic hydrodynamics [22], [2], [23], physically-informed neural network approaches [24], and data-driven learning approaches [25].

Compared to these previous methods that simulate underwater soft robots such as [22], our work models pneumatic actuation using the exact chamber geometry rather than artificial muscles facilitating greater accuracy at large deformations (see Fig. 1), uses a neural network thrust predictor rather than approximate analytical or heuristic

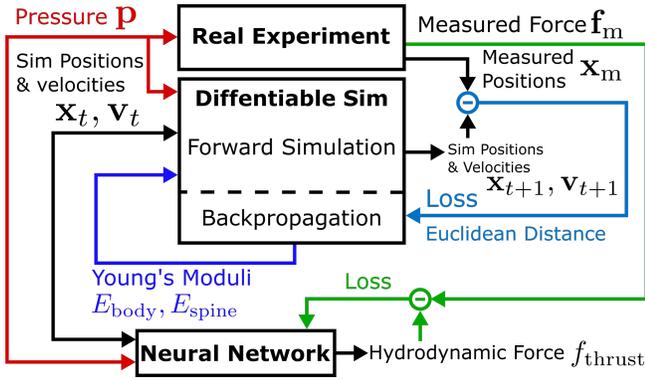


Fig. 2. Flow diagram of our differentiable simulation and learning pipeline. To learn the material parameters, we compute the loss as the Euclidean distance between the measured marker positions \mathbf{x}_m and the simulated position data \mathbf{x}_t and then minimize this loss using gradient-search enabled by our differentiable FEM simulation. The marker positions are 2D in-plane measurements primarily due to lateral displacement during flapping. Our neural network thrust predictor takes as input the positions and pressure at a previous time step and outputs the thrust force for the next time step.

hydrodynamics, and presents a more sophisticated hardware pipeline that can be used to validate simulation.

III. SYSTEM OVERVIEW

Our pipeline (Fig. 2) learns the material parameters of a soft robotic fish (see Fig. 1) using differentiable simulation and measurements from quasistatic experiments. Thrust force is simulated using a data-driven neural network predictor. We demonstrate that sim2real matching is good within 3% positional error normalized to actuator length for our soft robotic fish prototypes even under significant bending.

In Section IV-C, forward simulation is accomplished using implicit Euler time stepping and FEM spatial discretization. We leverage projective dynamics for a substantial speed up in computation [1]. We describe our bollard-pull style experimental setup in Section V. In Section VI-A, the material properties of the body and spine are specified by two different corotated materials stitched together at the boundary. We verify the results (Fig. 5) of our system identification by comparing simulations to measurements from a gamut of dynamic experiments at varying amplitudes and frequencies. Finally, in Section IV-B we describe the architecture and training of our thrust predictor network.

IV. PHYSICAL MODELING

A. Solid mechanics

The deformation of a soft material can be modeled by the equations of elasticity. For the constitutive model we choose an isotropic corotated material [26] since it is straightforward to implement for projective dynamics and can predict bending to within sufficient accuracy for engineering purposes.

A corotated material is defined by its energy density function using the Frobenius Norm and the trace operator,

$$\Psi = \mu \|\mathbf{U} - \mathbf{I}\|_F^2 + \frac{\lambda}{2} \text{tr}^2(\mathbf{U} - \mathbf{I}), \quad (1)$$

where \mathbf{U} is the stretch tensor, \mathbf{I} is the identity, and μ and λ are the Lamé parameters, which have the following relation to the Young's modulus E and Poisson's ratio ν for an isotropic material,

$$\mu = \frac{E}{2(1+\nu)}, \quad \lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}. \quad (2)$$

This material energy model is popular for physical simulation and animation, but it is not the most accurate choice for closing the sim2real gap for large deformations [26], [27].

To simulate a pneumatically actuated fish tail, we use a spatially-constant time-varying pressure boundary condition on the interior surface of each chamber of the tail.

B. Hydrodynamics

The hydrodynamic effects that produce thrust and drag on a swimmer are important, however, the physics governing the fluid can be complicated due to turbulent effects and two-way coupling between the solid and fluid domains. One may choose to model water using the Navier-Stokes equations, however, in practice simulating the full FSI problem is under time constraints computationally unfeasible. For this reason, many workers in the field of computer graphics and simulation choose to use simplified heuristic hydrodynamic models such as the one presented by Min et al. [23]. Analytical models for fish propulsion have also been studied extensively in the literature [28], [29], [30], but these approximations require severe simplifying assumptions necessary for analytical tractability. Using the elongated-body theory (EBT) described in Section IX-A [28], the average thrust produced by a fish is approximated by $f_{\text{EBT}} = m\hat{g}(\dot{x}, \dot{y}, y)$, where m is the virtual mass and \hat{g} is a function of the velocity \dot{x}, \dot{y} and deformation y of the fish tail measured in millimeters. We use EBT for comparison to our learned hydrodynamic model.

C. Simulation

1) *Discretization*: We use our finite element method with tetrahedra to discretize the continuous spatial domain occupied by the soft fish tail. Furthermore, we implement the implicit Euler time-stepping scheme for time discretization. Formally, let \mathbf{x}_i and \mathbf{v}_i be vectors of size $3N$ stacking up the nodal positions and velocities from the tetrahedra at the i -th time step. We then rewrite the governing equations in the following discretized form

$$\mathbf{x}_{i+1} = \mathbf{x}_i + h\mathbf{v}_{i+1}, \quad (3)$$

$$\mathbf{v}_{i+1} = \mathbf{v}_i + h\mathbf{M}^{-1}[\mathbf{f}_{\text{ela}}(\mathbf{x}_{i+1}) + \mathbf{f}_{\text{act}}]. \quad (4)$$

Here, h is the time step, M is the mass matrix, and \mathbf{f}_{ela} and \mathbf{f}_{act} represent the elastic force and the actuator force (computed by the pressure from the pneumatic actuator and the surface area of the chambers), respectively. The time integration is implicit because we make the elastic force depend on the nodal position \mathbf{x}_{i+1} at the beginning of the next time step instead of the current one.

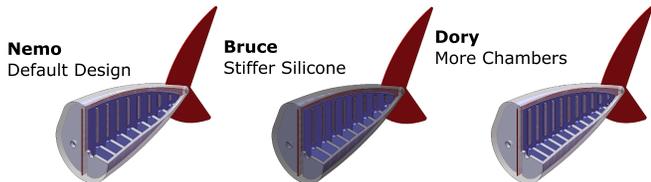


Fig. 3. Three different actuator designs. *Nemo* and *Bruce* share the same geometry but *Bruce* has a stiffer DS20 material for its body. *Nemo* and *Dory* share the same materials but *Dory* has a greater number of air chambers.

2) *Solver*: Solving Equation (3) typically requires expensive numerical computation due to its implicit time-stepping scheme. However, mature numerical tools exist since this problem has been extensively studied in computer graphics and physics simulation. In this work, we choose to use DiffPD [1], a recent differentiable simulator with projective dynamics that is suitable for simulating terrestrial and underwater soft robots. We choose differentiable projective dynamics over other numerical solvers, for example traditional Newton’s methods, due to its simultaneous accommodation of speed, robustness, and differentiability.

V. EXPERIMENTAL SETUP

In Fig. 3, we show renderings of the three types of fishtail actuators we investigate in this work. We vary material parameters as well as geometry.

We collect tail deformation, thrust force, and air chamber pressure data of pneumatically actuated silicone fishtail both in air and in water. To this end, the fishtail is mounted in a water tank and the fish head is rigidly connected to a load cell (TAL2210, 1 kg model) to measure the thrust force in the heading direction (Fig. 4). The load cell data is obtained using an amplifier board and a microcontroller.¹ Black markers are painted onto the fish’s back and tracked using Kanade-Lucas-Tomasi (KLT) feature tracking in a video captured from above using a GoPro Hero 6 camera [31], [32]. We found in practice that three markers are sufficient for capturing the deformation of the fishtail and can be used as input to our learning pipeline. From the video after correcting for parallax and distortion, we obtain in-plane displacements for each marker. An LED is used to synchronize the pressure and the force data with the recorded video. The fish is actuated with a Festo proportional valve manifold with controlled pressure², which outputs desired and actual pressure data for both air chamber sides of the fishtail.

Two types of experiments are conducted using the same experimental setup: First, a quasistatic experiment is carried out, where the fishtail is deformed by actuating one side of the fish with constant pressure. Second, flapping experiments are conducted where the fishtail is actuated using a square pressure wave, which leads to a natural motion of the tail. Both the quasistatic and the flapping experiment are carried out with a variety of fish designs using silicone rubber with

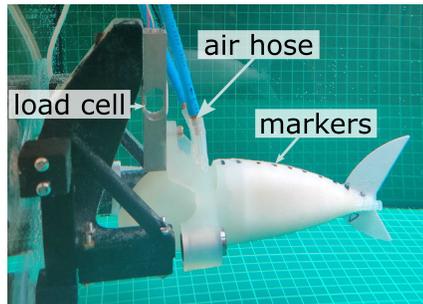


Fig. 4. Thrust experiment setup actuated with a Festo proportional valve manifold. A 3D-printed fixture (black) connects the load cell (grey) to a fish adapter (opaque). The fish adapter is mounted within a casted soft robotic fish tail made of silicone elastomer. The deformation of the tail is captured using a GoPro positioned above the tail. Markers on the back of the fish are tracked using Kanade-Lucas-Tomasi (KLT) feature tracking. The thrust generated by the fish is measured using a load cell.

TABLE I
MATERIAL, YOUNG’S MODULI, AIR CHAMBERS, AND MAXIMUM APPLIED PRESSURE FOR EACH PROTOTYPE. REPORTED MATERIAL PARAMETERS ARE NOMINAL VALUES.

| Property | <i>Nemo</i> | <i>Dory</i> | <i>Bruce</i> |
|-----------------------------|-------------|-------------|--------------|
| Body Material | DS10 | DS10 | DS20 |
| Body Young’s Modulus [MPa] | 0.1 – 0.25 | 0.1 – 0.25 | 1.1 |
| # of Chambers | 9 | 12 | 9 |
| Spine Young’s Modulus [GPa] | 2.5 – 5 | 2.5 – 5 | 2.5 – 5 |
| Max. Pressure [bar] | 0.2 | 0.35 | 0.5 |

different Young’s moduli.³ The fish materials, geometries, and actuation pressures are shown in Table I. No pre-straining was done on the fishtails prior to experiments and negligible hysteresis was found when actuating under periodic pressure signals. The maximum observed displacements in the fishtail were on the order of 20%, normalized to a tail length of 10 cm.

Since we are not fully constraining the fish to one-dimensional translation, there exists bending moments that may result in spurious forces on the load cell. However, we assume these to be negligible compared to the peak of the thrust force. Furthermore, due to the small size of the tank, a standing wave persists in the tank even after initial actuation. This can lead to a disturbance force on the load cell. However, since the frequency of this disturbance is known, it can be filtered out during post-processing. We also note that the measurement system has its own compliance and damping, which account for the recoil and negative force readings.

VI. SYSTEM IDENTIFICATION

A. Learning Material Parameters

In this section, we explain our method for system identification. The process begins with exporting two STL files from a CAD file of the soft actuator, one for the body and one for the spine. To streamline the tetrahedral mesh generation, we simplify the geometry by patching the small

¹Amplifier: HX711, Microcontroller: Arduino Leonardo

²MBA-FB-VI, 0 – 2 bar range, 1% accuracy

³Smooth-On Dragon Skin 10/20 slow with shore hardness 10A/20A

mounting holes used during the casting process to stabilize the fish. Otherwise, the geometry of the body and spine are imported as designed without modifications. In practice, we have found that a correct chamber and spine geometry are essential for adequate system identification of material parameters that are physically plausible.

We tetrahedralize the surface triangle mesh by modeling it as a whole piece of surface, a few holes representing the air chambers, and a set of internal points defining the separation of spine and body. The tetrahedralization is done according to the method by Hu et al. [33]. We control the target length of edges to be $1/50$ of the whole body length to get an expressive and representative tetrahedral mesh.

After the tetrahedralization, we split the spine and body elements and assign different Young’s moduli to them. The internal surfaces, which are identified as air chambers, are modeled as actuators to drive the tail. In physical experiments, we actuate the tail in the air and record its deformation through the tracked markers. This information provides supervision to the material parameter search for both the body and spine. We measure the error of matching by the Euclidean distance between measurement data and simulation results from the sensors.

We conduct our simulation experiments based on a linear elastic corotated material. The model can be improved by introducing a more sophisticated elastic energy model, for example Neo-Hookean material, which was not yet supported in our simulation framework. Furthermore, we chose not to include damping explicitly in the material model since we learn material parameters using static deformation data. For the dynamic experiments, we observed that the deformations were sufficiently small for accurate characterization and that the effect of material damping is small thus having a negligible effect on the forced response.

B. Learning Hydrodynamics

To learn a simplified model of the hydrodynamic thrust force predictor f_{thrust} , we use a feedforward neural network trained on data from the experimental setup (see appendix). The inputs of the neural network consist of the positions and velocities of the tracking points at time t , together with the measured pressure and its time derivative, while the output is a one-dimensional value for f_{thrust} at time t . The feedforward neural network consists of three hidden layers with RELU activation functions and respectively a dimension of 200, 300, and 200 hidden units. The network is trained with the Adam optimizer using a learning rate of 0.001 and a mean squared error loss. Note that our hydrodynamic modeling method does not estimate fluid properties such as density or viscosity, but rather the thrust force prediction, which is of more immediate use for designers of swimming robots.

In Table II, we provide a description of the training and the test set used for the thrust predictor network. Only *Nemo* and *Bruce* are reported in Fig. 8 and Table II for brevity since *Dory* thrust data predictions are similar. The table clarifies how we tested the generalizability of our predictor. We split the experimental data to show the ability of the network to

TABLE II

SPECIFICATION OF THE TRAINING AND TEST SETS FOR THE THRUST PREDICTOR NETWORK. IN THE NEMO EXPERIMENT, WE GENERALIZE TO UNSEEN TRIALS WITH SEEN PARAMETERS. IN THE BRUCE EXPERIMENT WE GENERALIZE TO UNSEEN HIGHER ACTUATION FREQUENCY.

| Name | Mode | N. trials | Pressure | Frequency |
|-------|------------|-----------|--------------|------------|
| Nemo | Training | 32 | 200,300 mbar | 1,2,3,4 Hz |
| | Validation | 8 | 200,300 mbar | 1,2,3,4 Hz |
| Bruce | Training | 30 | 500,750 mbar | 1,2,3 Hz |
| | Validation | 10 | 500,750 mbar | 4 Hz |

generalize to unseen trials. For the fishtail named *Nemo*, while we keep the parameters the same, we generalize in the validation set to unseen trials under various actuation amplitudes and frequencies. In the experiments for *Bruce*, we generalize to an unseen higher actuation frequency.

VII. RESULTS

A. Learned Material Parameters

For the *Nemo* tail actuator described in Table I, we perform a grid search of the loss landscape centered around the ground truth datasheet values for the silicone and acetal materials of the body and spine. We report that if the exact geometry of the actuator is reproduced with high fidelity in the simulation, we converge to values within the range of typical measured values for the material Young’s moduli (see Fig. 5). Further, we see that there is a unique minimum value that is within the acceptable range of measured moduli for both parameters. Note that we assume a priori that the Poisson ratio for silicone is $\nu \approx 0.5$ or *nearly* incompressible, a standard assumption for silicone, and the acetal sheet Poisson ratio is $\nu = 0.37$ as reported by the manufacturer. Typical values for the Young’s modulus of Dragon Skin 10 range from 0.1 MPa to 0.25 MPa and the value for Dragon Skin 20 was measured to be in the range of 1.1 MPa [34] and typical values for the Young’s modulus of acetal sheets range from 2.5 GPa to 5 GPa.⁴

In Fig. 6, we demonstrate that our method can be extended to higher dimensional parameter spaces such as a search over both Young’s moduli and Poisson’s ratios. Note that although we allow the Poisson ratio to vary in this identification experiment, the final value to which the body Poisson’s ratio converges is still *nearly* incompressible as expected for silicone.

B. Gradient-based and Gradient-free Solver Methods

We compare the runtime of the gradient-free method *CMA-ES* against the runtime of the gradient-based method *Adam* in Table III. The comparison shows that although *CMA-ES* [35] is slightly faster in runtime per iteration, the use of the gradient-based method *Adam* [36] is significantly more effective for convergence. These comparative experiments shown in Fig. 5 were carried out on a computer with an Intel Core i9-9900K @ 3.60GHz with 16 cores processor and 64.0 GB of memory.

⁴<https://dielectricmfg.com/knowledge-base/acetal/>

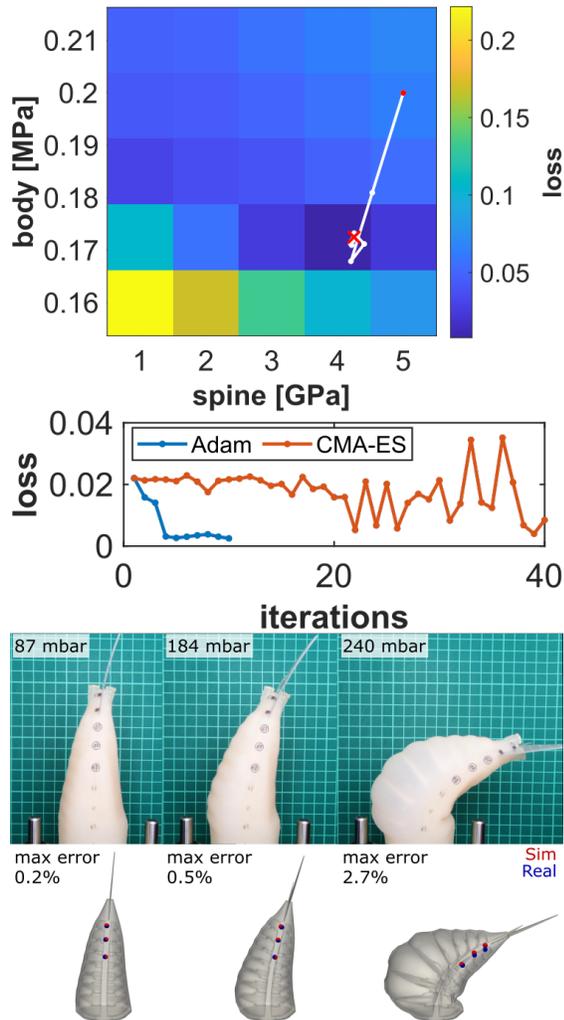


Fig. 5. *Top*: Exhaustive search of the Young’s moduli pair for the spine and body. The moduli pair with the lowest Euclidean loss is located at the red \times . The red dot indicates the start of a gradient search and the white line shows the progress towards convergence. *Middle*: Convergence comparison between a gradient-based (Adam) and gradient-free search (CMA-ES). *Bottom*: Comparison of static deformation of the *Nemo* fish for increasing pressure in experiment and simulation. Maximum displacement error increases with pressure, but remains within the measurement error (on the order of the diameter of markers). The grid in white under the real robot has 10 mm spacing. The reported error is normalized to the fish tail length of 10 cm.

C. Dynamic Experiments

For dynamic experiments, we compare the results of our simulation output with the measured data of the furthest tracked dot on the tail. In Fig. 7, we report our findings for sim2real performance in both the standard fish (*Nemo*) and two other data sets for a tail with different Young’s moduli for the body (*Bruce*) and a tail with a greater number of air chambers (*Dory*). We demonstrate that if system identification is done correctly our simulation results can predict the performance of a novel actuator design to within millimeter precision or 3% max normalized error using only a quasistatic data set for training without need for material testing. The *Bruce* prototype required higher pressures to get similar displacements to *Nemo* since the material is stiffer.

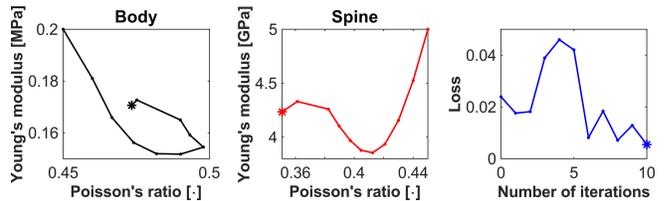


Fig. 6. Learning four material parameters from deformation data take with the *Nemo* fish prototype using a gradient-based approach that is run until convergence. The final values for Young’s moduli and Poisson’s ratios depicted as an asterisk agree with plausible material parameter values with lower loss.

TABLE III

COMPARISON OF ADAM, CMA-ES, AND GRID SEARCH. THE FORWARD SIMULATION TIME IS 318.9 SECONDS EQUIVALENT TO GRID SEARCH.

| Method | Total Iterations | Time Per Iteration | Total Time | Loss After 4 Iterations |
|-------------|------------------|--------------------|------------|-------------------------|
| Adam | 4 | 334.4 s | 22 m | 0.0027 |
| CMA-ES | 40 | 321.0 s | 214 m | 0.021 |
| Grid search | 25 | 318.9 s | 132.9 m | 0.023 |

The prototype *Dory* required higher pressure as well to produce similar displacements due to a greater number of actuation chambers. For videos of the dynamic experiments and simulation, we ask readers to refer to our supplemental video.

D. Learned Hydrodynamics

We compare our simple predictor of thrust force with the measurement data and the theoretical thrust from EBT, a classic, non-learning-based approximate analytical model from the literature, in Fig. 8. Although the model is capable of generalizing to actuation signals at frequencies previously unobserved in the training for the *Nemo* prototype, for the *Bruce* prototype the discrepancy is large, nearly twice the measured force, likely because of the more limited training data. In comparing to the thrust prediction from EBT, we see that the analytical thrust is strictly positive. The negative measured force is due to recoil of the measurement system. The advantage of our approach is the differentiability and flexibility of the neural network provided enough data is collected. We note that the network is capable of learning the measurement dynamics due to the compliance of the load cell and the damping of the water though these effects can be considered separately and corrected for if desired (see Section IX-B). We conjecture that the asymmetry of the measured thrust may be due to imperfections in the fabrication process of the fish tails favoring one direction more strongly.

VIII. CONCLUSION

We present an experimentally-verified simulation framework that can be used to accurately predict the deformations of a pneumatically actuated fish tail with a flexible spine. Our pipeline can accurately learn material parameters from a quasi-static data sets without having to do expensive and

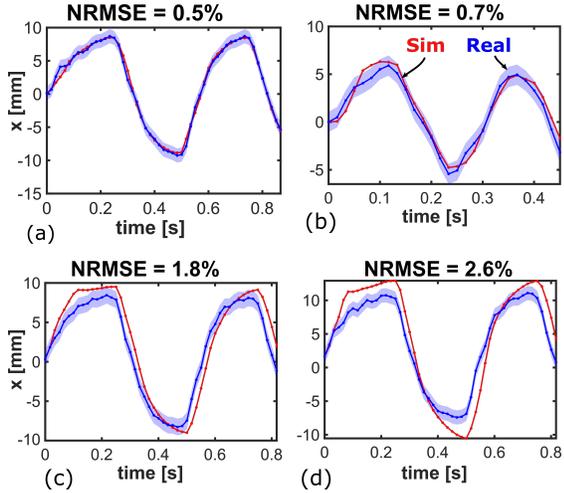


Fig. 7. Simulation and measurement data for the *Nemo* fish at (a) 200 mbar and 2 Hz, (b) 200 mbar and 4 Hz, (c) the *Bruce* fish at 500 mbar and 2 Hz, and (d) the *Dory* fish at 350 mbar and 2 Hz. The color bands indicate the variance from $N = 5$ trials. For both actuation signals, we are capable of achieving sub-millimeter accuracy between experiment and simulation. The same method is used to identify the parameters of Bruce and Dory, exhibiting accuracy still within 3 mm. We normalize the Root Mean Square Error (RMSE) to the fish tail length of 10 cm. As expected, higher pressures result in larger deformations with greater error. The phase lag exhibited in *Bruce* and *Dory* may be due to actuator dynamics present during higher actuation pressures.

time-consuming material testing. It also eliminates the need to do manual tuning of material constants to get accurate simulation results. The parameters we found are not only within typical range of measured material parameters for our materials, but can be used to successfully predict the behavior of dynamic experiments for different pressure actuation amplitudes and frequencies to within 3% positional error normalized to a actuator length of 10 cm. Although we use an isotropic corotated material, which is linear elastic, we find that this model is more sufficient to model large deformations on average giving acceptable displacement results for our engineering application. In these experiments, the damping of the material and the hydrodynamic effects are found to be negligible. This is because the actuation pressures used dominate the deformation compared to losses and hydrodynamic pressure.

We show a data-driven approach can be used to do simple prediction on a useful performance metric such as thrust force given a suitable hardware setup. However, more work is needed to produce a more robust thrust predictor if the morphology of the actuator changes substantially. We claim that for small design changes such as the choice of silicone or the number of internal chambers this framework can be used to quickly assess the relative merits of each design with a relatively sparse data set of approximately 30 types of different actuation signals.

Our aim is to further progress towards a systematic method by which soft roboticists can simulate and optimize their designs and controllers, whether they be soft fish, manipulators, or other flavors of soft robots. A fast and physically-verified co-optimization method of design and control is the goal.

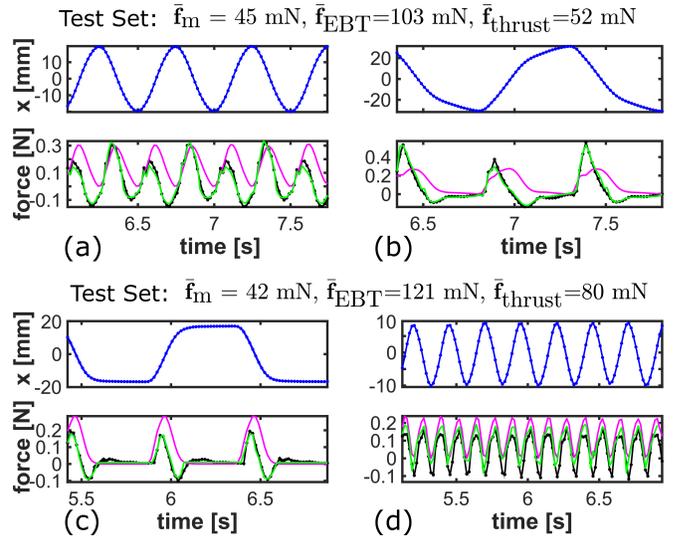


Fig. 8. Tail lateral displacement (blue) with thrust measurement and prediction. We compare the measured thrust force f_m (black), analytical thrust force from EBT f_{EBT} (magenta), and the neural network thrust prediction f_{thrust} (green). We show the thrust prediction for two training sets (a) and (c) and we report the time-averaged thrust for the two test cases (b) and (d). We note that EBT tends to over-predict the thrust measurement while our neural network thrust prediction accurately reproduces the frequency for a given actuation signal and matches amplitude more robustly than EBT.

IX. APPENDIX

A. Large Amplitude Elongated-Body Theory

The thrust of a fish swimming in an inviscid fluid can be described by Lighthill's Elongated-body theory (EBT) [28]. To calculate the thrust predicted by EBT, we make use of a simplified version of the reaction force expression from Lighthill [28], which uses significantly different notation,

$$f_{EBT} \approx \frac{1}{2} m \dot{x}^2, \quad (5)$$

where we have evaluated the expression for the marker closest to the tail and we have assumed the unit vectors $\frac{\partial x}{\partial a} \approx 0$ and $\frac{\partial y}{\partial a} \approx 1$ when considering small tail deformations in dynamic actuation. Note that \dot{x} is velocity tangential to the fin chord.

We also use Lighthill's expression for the virtual mass $m = \frac{1}{4} \pi \rho s^2$, where ρ is the density of the fluid and s is the cross-sectional depth, by estimating the geometry from the CAD model of our robotic fish.

B. Load Cell Measurement Dynamics

We choose to model the dynamics of the load cell measurement jig (Fig. 4) as a lumped element system shown in Fig. 9. We assume that the stiffness of the fixture is far greater than the stiffness of the connection between the load cell and the fish. Solving for the transfer function describing the load cell dynamics, we have

$$\frac{f_m}{f_{hydro}} = \frac{bs + k}{ms^2 + bs + k}, \quad (6)$$

where s is the complex Laplace variable. From this model, it is clear how the measured force f_m can be negative even for strictly positive f_{hydro} as discussed in Section VII-D.

It is possible to measure each parameter in Equation (6) and compensate it by inverting the transfer function, however, our proposed data-driven neural network force predictor is agnostic to the dynamics of the measurement system.

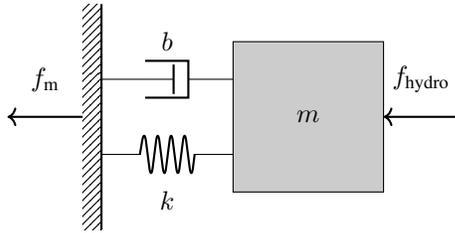


Fig. 9. Load cell dynamics. The force f_m is measured with a pre-calibrated load cell. The force f_{hydro} is the actual hydrodynamic force experienced by the fish. The stiffness k and the damping coefficient b describe the lumped mechanical impedance of the measurement system. The lumped mass m is dominated by the fish. The wall is the connection point to the load cell.

REFERENCES

- [1] T. Du, K. Wu, P. Ma, S. Wah, A. Spielberg, D. Rus, and W. Matusik, "Diffpd: Differentiable projective dynamics with contact," *arXiv preprint arXiv:2101.05917*, 2021.
- [2] P. Ma, T. Du, J. Z. Zhang, K. Wu, A. Spielberg, R. K. Katzschmann, and W. Matusik, "Diffaqua: A differentiable computational design pipeline for soft underwater swimmers with shape interpolation," *arXiv preprint arXiv:2104.00837*, 2021.
- [3] S. Bouaziz, S. Martin, T. Liu, L. Kavan, and M. Pauly, "Projective dynamics: Fusing constraint projections for fast simulation," *ACM transactions on graphics (TOG)*, vol. 33, no. 4, pp. 1–11, 2014.
- [4] Y. Hu, L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand, "DiffTaichi: Differentiable programming for physical simulation," *International Conference on Learning Representations*, 2020.
- [5] A. D. Marchese, R. K. Katzschmann, and D. Rus, "A recipe for soft fluidic elastomer robots," *Soft robotics*, vol. 2, no. 1, pp. 7–25, 2015.
- [6] R. K. Katzschmann, A. D. Marchese, and D. Rus, "Hydraulic autonomous soft robotic fish for 3d swimming," in *Experimental Robotics*. Springer, 2016, pp. 405–420.
- [7] R. K. Katzschmann, J. DelPreto, R. MacCurdy, and D. Rus, "Exploration of underwater life with an acoustically controlled soft robotic fish," *Science Robotics*, vol. 3, no. 16, 2018.
- [8] J. Zhu, C. White, D. K. Wainwright, V. Di Santo, G. V. Lauder, and H. Bart-Smith, "Tuna robotics: A high-frequency experimental platform exploring the performance space of swimming fishes," *Science Robotics*, vol. 4, no. 34, 2019.
- [9] G. Zheng, O. Goury, M. Thieffry, A. Kruszewski, and C. Duriez, "Controllability pre-verification of silicone soft robots based on finite-element method," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 7395–7400.
- [10] R. K. Katzschmann, M. Thieffry, O. Goury, A. Kruszewski, T.-M. Guerra, C. Duriez, and D. Rus, "Dynamically closed-loop controlled soft robotic arm using a reduced order finite element model with state observer," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2019, pp. 717–724.
- [11] M. Thieffry, A. Kruszewski, C. Duriez, and T.-M. Guerra, "Control design for soft robots based on reduced-order model," *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 25–32, 2018.
- [12] S. C. van den Berg, R. B. Scharff, Z. Rusák, and J. Wu, "Biomimetic design of a soft robotic fish for high speed locomotion," in *Conference on Biomimetic and Biohybrid Systems*. Springer, 2020, pp. 366–377.
- [13] Z. Wolf, A. Jusufi, D. Vogt, and G. Lauder, "Fish-like aquatic propulsion studied using a pneumatically-actuated soft-robotic model," *Bioinspiration & biomimetics*, vol. 15, no. 4, p. 046008, 2020.
- [14] M. Geilinger, D. Hahn, J. Zehnder, M. Bächer, B. Thomaszewski, and S. Coros, "Add: analytically differentiable dynamics for multi-body systems with frictional contact," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–15, 2020.
- [15] E. Coevoet, T. Morales-Bieze, F. Largilliere, Z. Zhang, M. Thieffry, M. Sanz-Lopez, B. Carrez, D. Marchal, O. Goury, J. Dequidt, *et al.*, "Software toolkit for modeling, simulation, and control of soft robots," *Advanced Robotics*, vol. 31, no. 22, pp. 1208–1224, 2017.
- [16] T. Du, K. Wu, A. Spielberg, W. Matusik, B. Zhu, and E. Sifakis, "Functional optimization of fluidic devices with differentiable stokes flow," *ACM Trans. Graph.*, vol. 39, no. 6, Nov. 2020. [Online]. Available: <https://doi.org/10.1145/3414685.3417795>
- [17] D. Hahn, P. Banzet, J. Bern, and S. Coros, "Real2sim: Visco-elastic parameter estimation from dynamic motion," *ACM Trans. Graph.*, vol. 38, no. 6, Nov. 2019.
- [18] Y. Hu, J. Liu, A. Spielberg, J. B. Tenenbaum, W. T. Freeman, J. Wu, D. Rus, and W. Matusik, "ChainQueen: A real-time differentiable physical simulator for soft robotics," *International Conference on Robotics and Automation*, 2019.
- [19] Z. Huang, Y. Hu, T. Du, S. Zhou, H. Su, J. Tenenbaum, and C. Gan, "Plasticinelab: A soft-body manipulation benchmark with differentiable physics," in *ICLR 2021*, 2021.
- [20] J. M. Bern, Y. Schnider, P. Banzet, N. Kumar, and S. Coros, "Soft robot control with a learned differentiable model," in *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2020, pp. 417–423.
- [21] M. Dubied, M. Y. Michelis, A. Spielberg, and R. K. Katzschmann, "Sim-to-real for soft robots using differentiable fem: Recipes for meshing, damping, and actuation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5015–5022, 2022.
- [22] T. Du, J. Hughes, S. Wah, W. Matusik, and D. Rus, "Underwater soft robot modeling and control with differentiable simulation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4994–5001, 2021.
- [23] S. Min, J. Won, S. Lee, J. Park, and J. Lee, "Softcon: simulation and control of soft-bodied animals with biomimetic actuators," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 6, pp. 1–12, 2019.
- [24] N. Wandel, M. Weinmann, and R. Klein, "Learning Incompressible Fluid Dynamics from Scratch – Towards Fast, Differentiable Fluid Models that Generalize," *arXiv:2006.08762 [cs, stat]*, Mar. 2021, arXiv: 2006.08762. [Online]. Available: <http://arxiv.org/abs/2006.08762>
- [25] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, "Neural ordinary differential equations," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 6572–6583.
- [26] E. Sifakis and J. Barbic, "Fem simulation of 3d deformable solids: a practitioner's guide to theory, discretization and model reduction," in *Acm siggraph 2012 courses*, 2012, pp. 1–50.
- [27] S. Adee, *Introduction to solid mechanics and finite element analysis using Mathematica*. Kendall Hunt, 2011.
- [28] M. J. Lighthill, "Large-amplitude elongated-body theory of fish locomotion," *Proceedings of the Royal Society of London. Series B. Biological Sciences*, vol. 179, no. 1055, pp. 125–138, 1971.
- [29] T. Y.-T. Wu, "Swimming of a waving plate," *Journal of Fluid Mechanics*, vol. 10, no. 3, pp. 321–344, 1961.
- [30] M. S. Triantafyllou, G. Triantafyllou, and D. Yue, "Hydrodynamics of fishlike swimming," *Annual review of fluid mechanics*, vol. 32, no. 1, pp. 33–53, 2000.
- [31] B. D. Lucas, T. Kanade, *et al.*, "An iterative image registration technique with an application to stereo vision." Vancouver, British Columbia, 1981.
- [32] C. Tomasi and T. K. Detection, "Tracking of point features," *Int. J. Comput. Vis.*, pp. 137–154, 1991.
- [33] Y. Hu, Q. Zhou, X. Gao, A. Jacobson, D. Zorin, and D. Panozzo, "Tetrahedral meshing in the wild," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 4, pp. 1–14, 2018.
- [34] L. Marechal, P. Bolland, L. Lindenroth, F. Petrou, C. Kontovounisios, and F. Bello, "Toward a common framework and database of materials for soft robotics," *Soft Robotics*, 2020.
- [35] C. Igel, N. Hansen, and S. Roth, "Covariance matrix adaptation for multi-objective optimization," *Evolutionary computation*, vol. 15, no. 1, pp. 1–28, 2007.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations*, 2015.