2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

A Data-Driven Approach to Geometric Modeling of Systems with Low-Bandwidth Actuator Dynamics

Siming Deng[1,2,*], Junning Liu[2], Bibekananda Datta[2], Aishwarya Pantula[3], David H. Gracias[3], Thao D. Nguyen[2], Brian A. Bittner[4,†], Noah J. Cowan[1,2,†,*]

***Abstract*— It is challenging to perform system identification on soft robots due to their underactuated, high-dimensional dynamics. In this work, we present a data-driven modeling framework, based on geometric mechanics (also known as gauge theory) that can be applied to systems with low-bandwidth control of the system's internal configuration. This method constructs a series of connected models comprising actuator and locomotor dynamics based on data points from stochastically perturbed, repeated behaviors. By deriving these connected models from general formulations of dissipative Lagrangian systems with symmetry, we offer a method that can be applied broadly to robots with first-order, low-pass actuator dynamics, including swelling-driven actuators used in hydrogel crawlers. These models accurately capture the dynamics of the system shape and body movements of a simplified swimming robot model. We further apply our approach to a stimulus-responsive hydrogel simulator that captures the complexity of chemo-mechanical interactions that drive shape changes in biomedically relevant micromachines. Finally, we propose an approach of numerically optimizing control signals by iteratively refining models, which is applied to optimize the input waveform for the hydrogel crawler. This transfer to realistic environments provides promise for applications in locomotor design and biomedical engineering.**

I. Introduction

Many conventional robots rely predominantly on rigid, fully actuated mechanisms. While these robots maintain superior force and precision compared to natural organisms, these rigid machines usually struggle in tasks that involve interacting safely with humans, handling deformable objects, and operating in unstructured environments [1]. Designs from nature have inspired the development of compliant mechanisms in robotics, enabling new capabilities [2]–[9]. The emergence of such soft components in robotic platforms has provided new avenues for improved adaptability, safety, cost, and energy efficiency. On the other hand, these new mechanisms introduce new challenges in modeling and control. The compliant nature of these soft components greatly increases the internal degrees of freedom as well as the degrees of underactuation, which makes it hard to obtain precise control of each shape element simultaneously. One potential approach is to utilize variations in passive actuation responses to stimuli across different parts of the body as a means of locomotion. This strategy capitalizes on the distinct temporal dynamics among subsystems or constituents.

In addition to compliance, underactuation, and low-bandwidth dynamics, there are other challenges introduced by soft robotic systems that complicate modeling efforts. For example, boundary conditions, surface interactions, and nonlinear material properties make it difficult to derive parsimonious models from physical principles. Using a top-down, data-driven approach, this work investigates the aforementioned passive responses within a systematic framework for locomotion control. We focus on systems with low-bandwidth shape changes in response to a single actuator input. Modeling actuator dynamics and its effects on the system can streamline engineering efforts to design and control soft robots, maximizing their capabilities with less exploratory or exhaustive experimentation. In many cases, a global model of system dynamics may be impractical to obtain, especially for custom machines, machines made through imprecise fabrication techniques, or systems deployed in poorly characterized environments (e.g. a patient's body). A key insight of this work is that in these scenarios, we can avoid reliance on a global model or the sample inefficiency of reinforcement learning schemes by systematically achieving function through iterative modeling and refinement of local behaviors.

Geometric mechanics provides a framework through which top-down insights permit novel motion planning approaches to dissipative systems with symmetry, including analytical computation of optimal gaits [10], [11]. A core premise of this work is that complex locomotor mechanics can be rewritten in a kinematic form, owing to the assumption that Rayleigh dissipation dominates the physical interaction between the body and environment [12]. The same framework has been instrumental in understanding cyclic locomotion in nature [13], [14]. Bittner et al. [15] presented a data-driven approach to construct a local model in the neighborhood of the observed limit cycle, using data points from stochastically perturbed, repeated behaviors. More recent work [16] extended this data-driven approach to shape-underactuated systems, which have high bandwidth control available only to a subset of the shape space. The ability to build local models provides the opportunity to

This material is based upon work supported by the National Science Foundation under grant No. #1830893, EFRI C3 SoRo: Programming Thermobiochemomechanical (TBCM) Multiplex Robot Gels

[1]Laboratory for Computational Sensing and Robotics, Johns Hopkins University, Baltimore MD 21218 USA.

[2]Department of Mechanical Engineering, Johns Hopkins University, Baltimore MD 21218 USA.

[3]Department of Chemical and Biomolecular Engineering, Johns Hopkins University, Baltimore MD 21218 USA.

[4]Johns Hopkins University Applied Physics Lab, Laurel, MD, 97331 USA.

[†]Co-supervised equally.

[*]To whom correspondence should be addressed: <sdeng10, ncowan>@jhu.edu

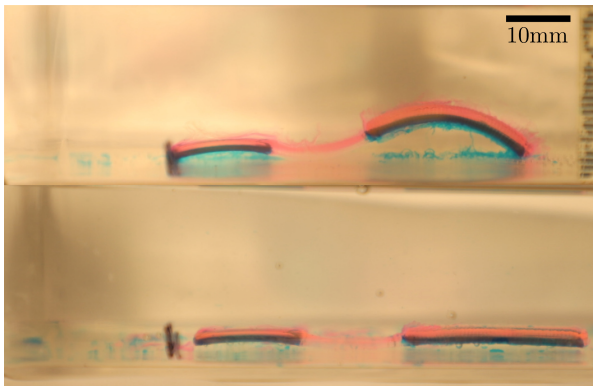

Fig. 1. Experimental screenshots of thermal cycling of two-segment robots with a flexible linker at the end of a cooling half-cycle (top), and heating half-cycle (bottom). One thermal cycle comprises a heating half-cycle and a cooling half-cycle. The robot displaces 4.4% body length at the end of the cycle.

sample candidate gaits offline for sample efficient hardware-in-the-loop optimization.

In contrast to hard robots that can readily be powered by multiple, independent electromagnetic modules with rapid (i.e., high-bandwidth) responses, soft-robots often must be stimulated by an external signal, typically a single stimulus such as pH, a specific biomolecule, light, or temperature. The impact of this signal must play out through complex kinetics (i.e., *low-bandwidth control*), where shape deformation is not instantaneously coupled to the control signal; rather, there is a temporal lag in the excitation of actuator dynamics following the control input. In our prior work [17] (see Fig. 1), we conceptualized and built a thermo-responsive hydrogel crawler. Although stimulus-responsive shape changes for hydrogels are ubiquitous in literature [18]–[20], the design of our robot exploited the swelling induced bending mechanism and asymmetry in response to the change in surrounding temperature to achieve the locomotion. In this crawler, there are three distinct segments: a suspended linker segment connects two end bilayer segments comprising active poly(*N*-isopropylacrylamide) (pNIPAM) and passive polyacrylamide (pAAM) layers with different morphologies. Asymmetry in friction forces, caused by morphological asymmetry, between the two bilayer segments at low and high temperatures allowed the robot to change its anchor during a temperature cycle to move unidirectionally. Additionally, we hypothesize that the distinct swelling rates of these bilayers induce a temporal asymmetry that can be exploited for locomotion.

Our core contribution, presented in Section III, is the data-driven modeling of a ubiquitous class of underactuated systems, where the shape dynamics are subject to a band-limited control. This contrasts with prior work [16], which required the application-limiting assumption that at least one element to be accessed by high bandwidth control. Our work enables the application of our framework for efficient behavior optimization and enhancement across a wide spectrum of previously unexplored soft robotic systems, such as hydrogel crawlers [17]. In IV, we demonstrate our methods on a well-known, analytically tractable system, modified to include low-bandwidth actuation of its shape parameters. Finally, in Section V we test our methods on a high-dimensional, finite element model of our previously published hydrogel robot [17]. In both examples, we show how the actuator dynamics can be simultaneously modeled with the body movements, enabling a data-driven modeling architecture for a broader class of soft or underactuated systems. Further, we use these examples to numerically optimize a parameterized input signal for certain objectives using an iterative parameter optimization and model refinement approach.

II. Background

A. Geometric locomotion model

Geometric mechanics [12], [21], [22] provides a framework for locomotion based on exploiting symmetry. Here, we consider a subclass of such group-invariant dynamical systems that are dominated by Rayleigh dissipation as caused by many types of isotropic friction [23]; in such dissipation-dominated systems, the equations of motion can be kinematically reduced such that the body velocity is expressed as a shape-dependent linear mapping of shape velocity. In this case, the kinematic equation is

$$\left(g^{-1}\dot{g}\right)^{\vee} = \xi = -\boldsymbol{A}(r)\dot{r}, \tag{1}$$

where ξ is the group velocity in its body frame[1], r denotes the system shape, and $\boldsymbol{A}(r)$ is called the *local connection*. Here the matrix $\boldsymbol{A}(r)$ is a function of shape r and acts analogously to a Jacobian in which it relates the system's shape velocities to body velocities. A spatial trajectory of the system body frame can be calculated by integrating (1) with respect to a fixed reference frame.

B. Data-driven modeling

Bittner et al. [15] developed a data-driven approach to geometric modeling and optimization, which was later extended [16] to be applied to systems with high bandwidth control in only a subset of the shape variables. This approach allows local estimation of a connection in the neighborhood of a limit cycle with far fewer samples than required to train a global model, making it practical for *in-situ* system identification, especially for systems with high dimensional shape spaces.

In this approach, system shape trajectories are fit to an oscillator such that each data point is assigned a phase value [24]. Then, a local Taylor approximation of the connection can be computed via linear regression across data points within phase windows.

We detail the process by which we estimate a linearized model within each phase window. Data-driven Floquet analysis techniques extract information from the observed oscillator data and assign each sample point an estimated phase [24], [25]. The observed shape samples are then phase-averaged and fitted to a Fourier series to obtain a limit cycle,

[1]Here $(\cdot)^{\vee}$ is an isomorphism that maps velocities from Lie algebra form to vector form, and its inverse is denoted as $\widehat{(\cdot)}$. In a SE(2), $(\cdot)^{\vee} : \mathrm{se}(2) \to \mathbb{R}^3$, and $\widehat{(\cdot)} : \mathbb{R}^3 \to \mathrm{se}(2)$

denoted as $\theta_r(\cdot)$. The perturbed trajectory, relative to the limit cycle, is denoted as $\delta_r := r - \theta_r$. The first-order Taylor approximation of the local connection in each local phase window can be constructed as

$$\boldsymbol{A}_k(r) \approx \boldsymbol{A}_k(\theta_r) + \delta_r^T \frac{\partial \boldsymbol{A}_k}{\partial r}, \tag{2}$$

where $\boldsymbol{A}_k(r)$ is the k^{th} row of the local connection, which is a vector of the same dimension as shape perturbation δ_r.

All samples are grouped into neighborhoods by their estimated phase values, and a local model is fitted in each phase window. In the m^{th} phase window, the averaged shape is assumed to be a constant θ_r^m. The first-order Taylor approximation of the local connection matrix $\boldsymbol{A}(r)$ in this phase window can be fitted by solving the following Generalized Linear Model (GLM):

$$\xi_k^{(n)} \sim \boldsymbol{C}_k + \boldsymbol{B}_k \delta_r^{(n)} + \boldsymbol{A}_k(\theta_r)\dot{\delta}_r^{(n)} + \frac{\partial \boldsymbol{A}_k}{\partial r}\delta_r^{(n)}\dot{\delta}_r^{(n)}. \tag{3}$$

Here, $\xi_k^{(n)}$ corresponds to the k^{th} coordinate of the n^{th} sampled body velocity $\xi^{(n)}$, and $\delta_r^{(n)} := r^{(n)} - \theta_r^m, \dot{\delta}_r^{(n)} := \dot{r}^{(n)} - \dot{\theta}_r^m$ are the shape and shape velocity perturbation samples defined in the local region indexed by m. Regressor $C_k := \boldsymbol{A}_k(\theta_r)\dot{\theta}_r$ describes the average behavior in the neighborhood of θ_r^m. $B_k := \dot{\theta}_r^T \frac{\partial \boldsymbol{A}_k}{\partial r}$ and $\boldsymbol{A}_k$ are the terms that respectively relate the effects of shape and shape velocity offsets from the limit cycle. $\frac{\partial \boldsymbol{A}_k}{\partial r}$ is the cross term that incorporates the interaction between δ_r and $\dot{\delta}_r$. Note that this is a local estimate in the m^{th} phase window. This local approximation is repeated for all separate groups of data points, after which a Fourier series model is used to guarantee a smooth transition among the fitted matrices.

III. Methods

A. Low-bandwidth shape control

In this paper, we consider systems whose locomotion can be characterized by (1) while only having access to low-bandwidth control over r. In particular, we assumed the dynamics on r to take the general form of

$$\dot{r} = f(r, u), \tag{4}$$

where the system shape velocity $\dot{r}$ is a function of its shape r and an input u.

First, we extracted a phase-averaged gait cycle (θ_r, θ_u) from the general input [24]. Denoting the perturbation from phase-averaged shape and control as $\delta_r := r - \theta_r, \delta_u := u - \theta_u$, the local first-order Taylor approximation of the actuation dynamics can be written in the following form:

$$f(r, u) \approx f(\theta_r, \theta_u) + \frac{\partial f}{\partial r}(\theta_r, \theta_u)\delta_r + \frac{\partial f}{\partial u}(\theta_r, \theta_u)\delta_u \tag{5}$$

We then fit the data to the above first-order approximation by solving the following Generalized Linear Model,

$$\dot{\delta}_r^{(n)} \sim \boldsymbol{D} + \boldsymbol{E}_r \delta_r^{(n)} + \boldsymbol{E}_u \delta_u^{(n)}, \tag{6}$$

where D is the average shape velocities of the observed data in the local phase window, and $(\boldsymbol{E}_r, \boldsymbol{E}_u)$ are the terms that describe how shape and input offsets respectively modify the average behavior. $\delta_r^{(n)} := r^{(n)} - \theta_r^m, \delta_u^{(n)} := u^{(n)} - \theta_u^m$ are the shape and input perturbations defined in the m^{th} local phase window, where (θ_r^m, θ_u^m) are the mean values of shape and input.

The fitted models from (6) and (3) can be used in series to make predictions of the system shape and position trajectories given the input signal. First, the input signal $u(t)$ is transformed into the phase coordinate using the fitted phase map. The initial shape is assumed to be on the limit cycle ($\delta_r = 0$) as the same phase value of the initial input $u(t_0)$. At each discrete time t_i, the shape velocity perturbation $\dot{\delta}_r(t_i)$ is predicted using the actuator model (6) given the current shape perturbation $\delta_r(t_i)$ and the input perturbation $\delta_u(t_i)$. $\delta_r(t_i)$ is then integrated by Euler's method to obtain the predicted shape at the next time step $\delta_r(t_{i+1})$. The predicted shape $\delta_r(t_i)$ is then used to predict the body velocity $\xi(t_i)$ using the body velocity model (3). The predicted body frame position at the next time step $g(t_{i+1})$ is then integrated using $\xi(t_i)$; the two integration steps of each model evolve in series.

We start with knowledge of the initial system shape $r(t = 0)$ and the control input $u(t)$. Then we can numerically solve (4) and (1) together using the fitted regression models, (6) and (3).

We apply the model improvement metric described in [16], comparing our first-order regression model predictions to the phase-averaged baseline model predictions,

$$\Gamma_\chi = 1 - \frac{\sum_{n=1}^{\mathcal{N}} \|\chi_D^{(n)} - \chi^{(n)}\|}{\sum_{n=1}^{\mathcal{N}} \|\chi_T^{(n)} - \chi^{(n)}\|}. \tag{7}$$

This improvement metric is defined as one minus the relative error of the data-driven prediction χ_D with respect to the baseline prediction χ_T over $\mathcal{N}$ samples of body velocity and shape velocity $\chi = \{\xi, \dot{r}\}$. $\Gamma_\chi \leq 0$ means the data-driven prediction is no better than the phase-averaged prediction, and $0 < \Gamma_\chi \leq 1$ means that our model can make better predictions than the baseline model, up to perfect reconstruction of the ground truth at $\Gamma_\chi = 1$.

The baseline phase model corresponds to the zeroth-order model, generating predictions solely reliant on current phase information. This metric holds significance in assessing the extent to which our data-driven first-order model outperforms the baseline model within the local region. A substantial improvement metric implies that the data-driven model more accurately approximates the ground truth compared to the baseline model within the local region, making it more suitable for local optimization.

B. Optimizing behaviors and iterative model refinement

The expense of data collection[2] incentivizes our focus on sample-efficient optimization schemes. Given an input parameterization, we sparsely sample data in the full range of the input space and build a rough model, which was used to locally optimize the input parameters. Then we zoom into the region around the optimized parameter and re-sample points in this local area. The model built with sample points

in a smaller region will be more localized and accurate over that confined domain. We iterate between these two processes—optimization and model refinement—so that in the end it converges to a local optimum in the control space. A global optimum is not guaranteed. In Sections IV and V, we demonstrate the methods using a sample objective where we maximize the displacement per cycle and penalize the cycle time.

IV. Illustrative example: Purcell swimmer with low-bandwidth actuation

Before demonstrating the method on data, it is helpful to interrogate a simple analytical model, such as a three-link Purcell swimmer. We modify the model to include low-bandwidth actuation, inspired by the low-bandwidth actuation of our hydrogel robot:

$$\dot{r_i} = c_i(r_i^s(T) - r_i), \quad c_i > 0, i = 1, 2, \tag{8}$$

where r_i is the i^{th} shape variable (joint positions), $r_i^s(T)$ denotes each joint's steady state equilibrium given temperature, and c_i is the converging rate of each joint towards its steady state equilibrium. Specifically, the steady-state equilibrium $r_i^s(T)$ is assigned to be a linear function of a one-dimensional input signal, temperature T. We have a bound on temperature that puts limits on the swimmer's joint angles. The resulting shape trajectory exhibits hysteresis which is often observed in low-bandwidth control systems.

Assuming different constants c_i on the two joints, the shape variables will exhibit a gait where both joints are not synchronized under a repeating temperature cycle. Although both joints are controlled by the same temperature input, the phase lag between the two joints breaks the symmetry of joint synchronization, making the gait enclose a nonzero area in the shape space, which is critical for locomotion in viscous swimming domains from the scallop theorem [26]. The greater the difference in c_i, the larger the phase lag is.

A. Input generation

Our parameterization on the control signal is concise while maintaining the ability to alter important features of the temperature profile. Here, we used 4 parameters to describe the temperature cycle: a low-point temperature T_{low}, a high-point temperature T_{high}, time-span per cycle t_{cycle}, and the portion of the half period to ramp between high and low temperatures $\eta_{\text{ramp}} = 2t_{\text{ramp}}/t_{\text{cycle}}$, where t_{ramp} is the time to ramp between high and low temperatures. Performing multiple cycles of these parameterized temperature cycles, the shape trajectory forms a stable orbit under periodic forcing. We then perturbed the forcing parameters across cycles, which resulted in what can be seen as a "tube" around the orbit as shown in Fig. 3.

[2]In practice, a typical temperature cycle for our hydrogel crawler takes 6 hours because of the slow actuation kinetics of the material. The FEA simulation is computationally expensive as well (given the minimum mesh size that produces fair simulations); running a 10-cycle simulation on a well-equipped desktop computer takes about 2 days. Both facts make data collection for such systems expensive, thus data efficiency is crucial.

V. Main application: hydrogel crawler

A. Actuator dynamics

Bilayers and other multi-material structures are useful in creating interesting modes of shape changes like bending. Typical swelling-driven bilayer bending dynamics are similar to the form of an exponential low-pass filter mentioned in IV. Specifically, the geometry of the bilayer (e.g., layer thickness ratio and material properties) can affect the steady state equilibrium of the shape variables as well as the rate of reaching equilibrium.

B. Hydrogel crawler

Thermo-responsive hydrogel crawlers in [17], capable of swelling and shrinking, utilize geometric asymmetry, leading to asymmetry in friction forces, to generate net motion under temperature cycles. We utilized the same 2D finite element model in Abaqus Unified FEA [27] to produce a time-dependent $x - y$ coordinate along the contour of the robot to calculate the area (2-D volume) of the actuated segments. The data is then parameterized into the shape variable r.

Here we assume the actuation dynamics in a general (nonlinear) form (4), without any specific structure on it, namely

$$\dot{r} = f(r, T), \tag{9}$$

where the input is assumed to be the temperature T.

C. Finite element model

Briefly, our finite element model, based on chemo-mechanics described in [28], solves coupled diffusion-deformation equations for hydrogel undergoing temperature-driven swelling and shrinking. We used Neo-Hookean and Flory-Huggins potentials to describe the entropic elastic behavior of the polymer network and the mixing of polymer-solvent, respectively. The swelling of pNIPAM caused by the lower critical solution temperature transition (LCST) was modeled by assuming a sigmoidal function for the temperature dependence of the Flory-Huggins interaction parameter. We also assumed that the diffusivity of the water through pNIPAM also increased sigmoidally with temperature across the LCST, which caused the characteristic time of deswelling to be significantly faster than the characteristic time of swelling. Given the relatively long actuation time, the environment temperature can be simply assumed to be evenly distributed. We also considered a combined effect of gravity and buoyancy by prescribing a net body force on the hydrogel structure. Our material model parameters were either directly determined from experiments or calibrated against experiments using finite element analysis. Further details of the finite element simulation can be found in [17].

D. Input generation

Here we used a similar parameterization for the input temperature signal as that in IV-A. However, for thermo-responsive hydrogels, the kinetics of the swelling and shrinking vary distinctively. We, therefore, separated the input cycle into two independent parts, cooling and heating. Thus, the

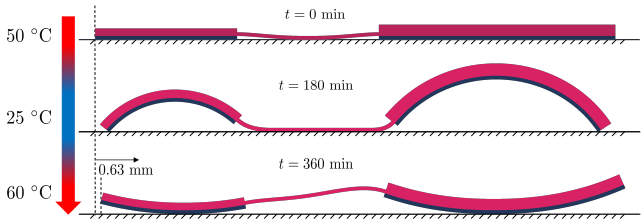

Fig. 2. A representative finite element analysis showing unidirectional motion towards the larger bilayer of the thermo-responsive hydrogel crawler subjected to 3 hours of ramped cooling and 3 hours of ramped heating cycle.

Parameter	Range
T_{low}, T_{high}	$20 \sim 41$, $45 \sim 65$ $°C$
t_{cool}, t_{heat}	$2 \sim 8$, $0.5 \sim 3$ hrs
η_{cool}, η_{heat}	$\frac{1}{32} \sim 1$, $\frac{1}{32} \sim 1$

TABLE I

FULL INPUT PARAMETER RANGES.

dimension of the input parameter space increases to six, low temperature T_{low}, ramped cooling time span t_{cool}, cooling ramp time ratio η_{cool}, high temperature T_{high}, ramped heating time span t_{heat}, and heating ramp time ratio η_{heat}. The ramp time ratios are defined as the ratio of the corresponding ramp time to the cooling or heating time span, i.e., $\eta_{\text{cool}} = t_{\{\text{ramp,cool}\}}/t_{\text{cool}}$. The allowable range of each parameter is determined by material properties and characteristic diffusion time and was validated by a parametric study using FEA. Specifically, the ramped cooling and heating timespans are determined by scaling swelling and shrinking characteristic time of the hydrogel, both temperature ranges are specified by 4% equilibrium strain span of the material. The ramp ratios are in the range of $[0, 1]$[3]. The calculated full input parameter ranges are shown in Table I. The noisy input signal is generated by sampling from a uniform distribution in the parameter space. At each iteration, we ran a total of 100 cycles of FEA simulation.

E. Shape parameterization

The compliant nature of the devices and external forces makes the internal shape high-dimensional. However, fitting models to a highly dimensional shape space will likely cause overfitting. We thus seek a reduced-order representation of the shape of the system. Here, principal components analysis (PCA) is a simple candidate reduction method that could serve this purpose. While it is a straightforward way of reducing the effective degrees of freedom, the complex coupling between segments led to principal components that lacked a clear physical interpretation. As an alternative, we considered that the volume of each active section is a more physical, descriptive candidate for representing the system shape variables. This parameterization provided a clear, intuitive relationship between the two segments, and exhibited the phase lag between the smaller and larger bilayer segments that we expected from the design.

[3]Ramp ratio near 0 results in temperature jumps, which is impractical and often causes stability and convergence issues in FEA simulation because of the excessive deformation of the finite element mesh in a short period of time. Thus in the implementations, we raised the lower bounds to $\frac{1}{32}$.

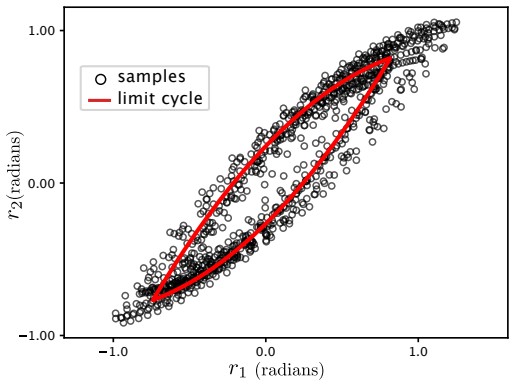

Fig. 3. Sampled shape data points around a limit cycle. Note that we cannot directly sample the full shape and shape velocity space due to the presence of the actuator dynamics. Fortunately, adding perturbation to the input parameters does generate shape and shape velocity variations around the limit cycle, which is essential for building models of behaviors.

F. Input optimization

As a demonstration, we optimized the input parameters to maximize the displacement per cycle. The objective function is defined as

$$F(u) = \Delta g_x - \lambda t_{\text{cycle}}, \tag{10}$$

where Δg_x is the displacement in the x direction per cycle, and t_{cycle} is the cycle time. λ is a penalty factor that controls the trade-off between the above two terms. The objective function is to maximize the net displacement per cycle. During the optimization process, we noticed that the optimizer tended to find cycles with the longest possible cycle times, optimizing cycle-to-cycle distance, with no penalty on time, pushing the results toward the boundary. To address this, we added a regularizing term to penalize the cycle time.

We started by sampling 100 points (100 cycles of system motion) in the full input parameter space (see Table I). We performed ten-fold cross-validation to avoid overfitting. A rough model of the system was built using the initially sampled data points, and then the model was used to optimize for an input parameter that maximizes the objective function above. The optimization was performed using the Sequential Least Squares Programming algorithm where the local gradient and Hessian were estimated using finite differencing. Once the numerical optimum was obtained, we shrank the range of each input parameter by 35%, centering at the optimum, and repeated the optimization process. We repeated this process three times, and the parameters converged to a performant gait. The model improvement metrics were calculated for each iteration as shown in Fig. 4

VI. Discussion and Conclusion

In this work, we designed and implemented a data-driven modeling framework for dissipative systems with low-bandwidth actuator dynamics. We showed the success of this method in predicting behaviors on a classical simplified model, the Purcell swimmer, with a modified class of passive shape dynamics. We built on prior work, relaxing the requirement that at least one shape element is accessed through high-bandwidth control. In doing so, our

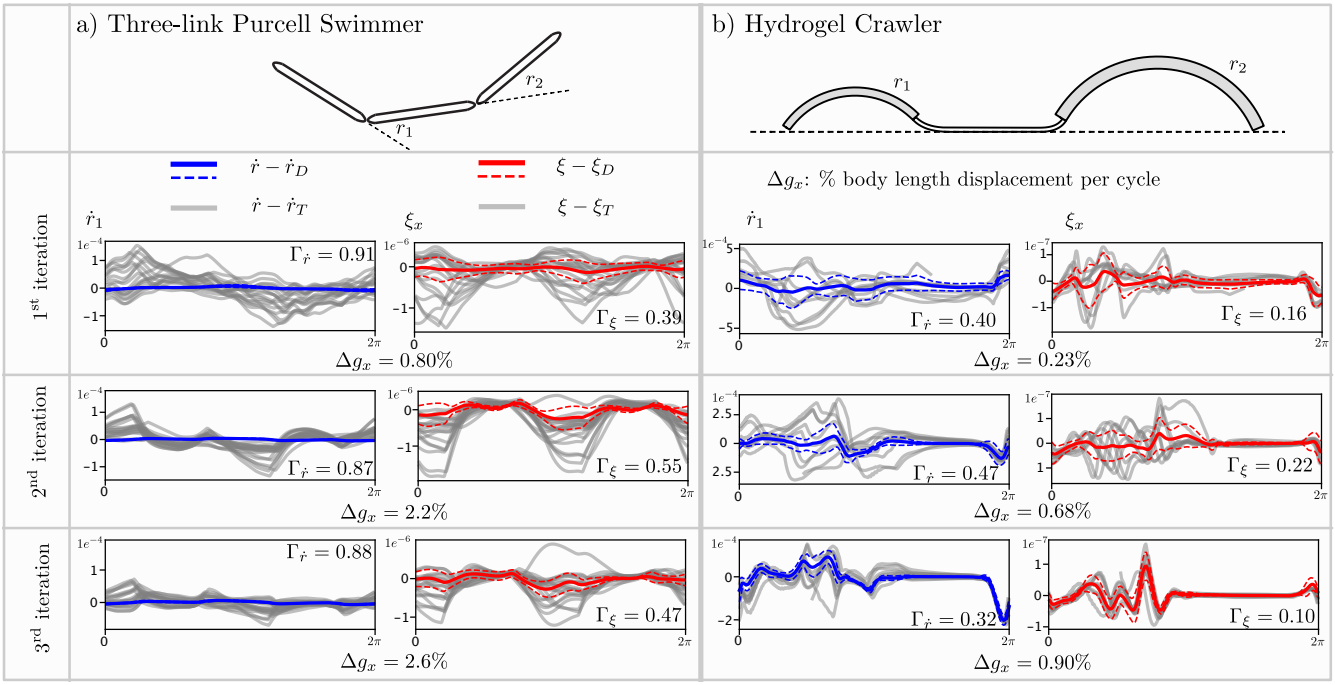

Fig. 4. Model prediction metrics for the Purcell swimmer (top left) and hydrogel crawler (top right) model. We plot sample model prediction errors in phase coordinates from each of the model refining iterations. Also, the average displacement, in terms of percentage body length per actuation cycle, each iteration is shown below the plots. The metrics are calculated on the testing set which is unseen by the model both on shape velocity predictions and body velocity predictions. Test set trajectories (grey) along with the mean (solid) and standard deviation (dashed) of the data-driven model prediction error are plotted for each iteration. The body velocity prediction error is inherently propagated from the shape velocity prediction error, and thus the body velocity prediction errors are unsurprisingly greater than the shape velocity prediction error. While the model prediction errors start to drop as we shrink the sampling range due to overfitting, the average behavioral objective improves each iteration.

method enables the modeling of novel mechanisms like the hydrogel crawler, whose internal degrees of freedom all exhibit a passive response to controllable stimuli. Despite the challenges often associated with designing control signals for such actuators, the robot was intentionally designed to capitalize on morphology-induced actuation asymmetries for locomotion. We showed not only that we could model the crawler with accuracy beyond the phase-averaged gait, but that the system was capable of using this model in a gradient-based optimization scheme to rapidly identify a viable crawling maneuver. The broader implications of this result are that we now have a rational framework to pursue data-driven modeling and optimization of a much larger class of underactuated systems. For applications in biology, where continuous, soft interfaces facilitate safe interaction with the body, this method provides the potential to model new mechanisms pre-deployment in the body and even *in situ*, since variation amongst morphology and environment across patients can be significant. Key additional future efforts will include power, actuation, and sensing at the scales desired for the locomotion application. In addition, it is well known that there is a significant phase lag between muscle activation and body movements [29], suggesting that our approach can be used to better understand the underlying neural control problem [30].

The model improvement saturated and decayed as the sampling region was reduced in the hydrogel crawler gait optimization result. This likely occurred because there was less variation in the sampled data (our tube of gait distribution data had a smaller overall volume). This is loosely analogous to convergence results in adaptive control, which often rely on sufficient excitation of the dynamics; in a similar vein, we do not expect to learn informative improvement without cycles that excite significant dynamic variation. We have shown that the model was iteratively improved and the optimizer settles at a performant gait.

In fact, many next-generation soft robots will be hard to model apriori. Imagine custom morphologies deployed in uncharacterized patient bodies (such as an artificial heart valve [31]). This mechanism must leverage the ability to build a viable control policy from limited data. We have provided a framework through which it and similar robots can rapidly obtain functional motion primitives. We juxtapose our method with reinforcement learning [32], deep learning architectures [33], and gait optimization [34] which require large amounts of data that are unavailable in simulation and impractical to empirically obtain for many applications (e.g. in biomedical settings).

While many samples might be available in simulated environments, there are many platforms that must be system-identified in the field. *In-situ* system identification (such as the type we implement here) paired with a gradient-based optimizer provides a tractable, system-oriented way to pursue optimization of robot behaviors in hard-to-model environments. Further efforts can be made to wrap control, modeling and optimization into an in-the-loop process, where the robot can learn to navigate in complex, time-varying environments without human intervention.

REFERENCES

[1] D. Trivedi, C. D. Rahn, W. M. Kier, and I. D. Walker, "Soft robotics: Biological inspiration, state of the art, and future research," *Applied Bionics and Biomechanics*, vol. 5, pp. 99–117, 2008.

[2] R. F. Shepherd, F. Ilievski, W. Choi, S. A. Morin, A. A. Stokes, A. D. Mazzeo, X. Chen, M. Wang, and G. M. Whitesides, "Multigait soft robot," *Proceedings of the National Academy of Sciences*, vol. 108, no. 51, pp. 20 400–20 403, 2011.

[3] S. Li, H. Bai, Z. Liu, X. Zhang, C. Huang, L. W. Wiesner, M. Silberstein, and R. F. Shepherd, "Digital light processing of liquid crystal elastomers for self-sensing artificial muscles," *Science Advances*, vol. 7, no. 30, p. eabg3677, 2021.

[4] T. TolleyMichael, F. ShepherdRobert, C. GallowayKevin, J. WoodRobert, M. WhitesidesGeorge, *et al.*, "A resilient, untethered soft robot," *Soft Robotics*, 2014.

[5] B. Liu, Y. Ozkan-Aydin, D. I. Goldman, and F. L. Hammond, "Kirigami skin improves soft earthworm robot anchoring and locomotion under cohesive soil," in *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2019, pp. 828–833.

[6] Y. O. Aydin, J. L. Molnar, D. I. Goldman, and F. L. Hammond, "Design of a soft robophysical earthworm model," in *2018 IEEE International Conference on Soft Robotics (RoboSoft)*. IEEE, 2018, pp. 83–87.

[7] H. Zhao, Y. Li, A. Elsamadisi, and R. Shepherd, "Scalable manufacturing of high force wearable soft actuators," *Extreme Mechanics Letters*, vol. 3, pp. 89–104, 2015.

[8] H. Lu, M. Zhang, Y. Yang, Q. Huang, T. Fukuda, Z. Wang, and Y. Shen, "A bioinspired multilegged soft millirobot that functions in both dry and wet conditions," *Nature Communications*, vol. 9, no. 1, p. 3944, 2018.

[9] W.-H. Chen, W. Yang, L. Peach, D. E. Koditschek, and C. R. Sung, "Kinegami: Algorithmic design of compliant kinematic chains from tubular origami," *IEEE Transactions on Robotics*, vol. 39, pp. 1260–1280, 2022.

[10] R. L. Hatton and H. Choset, "Nonconservativity and noncommutativity in locomotion," *The European Physical Journal Special Topics*, vol. 224, no. 17, pp. 3141–3174, 2015.

[11] R. L. Hatton, T. Dear, and H. Choset, "Kinematic cartography and the efficiency of viscous swimming," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 523–535, 2017.

[12] S. D. Kelly and R. M. Murray, "Geometric phases and robotic locomotion," *Journal of Robotic Systems*, vol. 12, no. 6, pp. 417–431, 1995.

[13] Y. Ozkan Aydin, B. Chong, C. Gong, J. M. Rieser, J. W. Rankin, K. Michel, A. G. Nicieza, J. Hutchinson, H. Choset, and D. I. Goldman, "Geometric mechanics applied to tetrapod locomotion on granular media," in *Biomimetic and Biohybrid Systems: 6th International Conference, Living Machines 2017, Stanford, CA, USA, July 26–28, 2017, Proceedings 6*. Springer, 2017, pp. 595–603.

[14] D. Zhao, B. Bittner, G. Clifton, N. Gravish, and S. Revzen, "Walking is like slithering: A unifying, data-driven view of locomotion," *Proceedings of the National Academy of Sciences*, vol. 119, no. 37, p. e2113222119, 2022.

[15] B. Bittner, R. L. Hatton, and S. Revzen, "Geometrically optimal gaits: a data-driven approach," *Nonlinear Dynamics*, vol. 94, no. 3, pp. 1933–1948, 2018.

[16] ——, "Data-driven geometric system identification for shape-underactuated dissipative systems," *Bioinspiration & Biomimetics*, vol. 17, no. 2, p. 026004, jan 2022.

[17] A. Pantula, B. Datta, Y. Shi, M. Wang, J. Liu, S. Deng, N. J. Cowan, T. D. Nguyen, and D. H. Gracias, "Untethered unidirectionally crawling gels driven by asymmetry in contact forces," *Science Robotics*, vol. 7, no. 73, p. eadd2903, 2022.

[18] J. C. Breger, C. Yoon, R. Xiao, H. R. Kwag, M. O. Wang, J. P. Fisher, T. D. Nguyen, and D. H. Gracias, "Self-folding thermo-magnetically responsive soft microgrippers," *ACS applied materials & interfaces*, vol. 7, no. 5, pp. 3398–3405, 2015.

[19] D. Han, C. Farino, C. Yang, T. Scott, D. Browe, W. Choi, J. W. Freeman, and H. Lee, "Soft robotic manipulation and locomotion with a 3d printed electroactive hydrogel," *ACS applied materials & interfaces*, vol. 10, no. 21, pp. 17 512–17 518, 2018.

[20] H. Na, Y.-W. Kang, C. S. Park, S. Jung, H.-Y. Kim, and J.-Y. Sun, "Hydrogel-based strong and fast actuators by electroosmotic turgor pressure," *Science*, vol. 376, no. 6590, pp. 301–307, 2022.

[21] H. Cendra, J. E. Marsden, and T. S. Ratiu, "Geometric mechanics, lagrangian reduction, and nonholonomic systems," *Mathematics Unlimited—2001 and Beyond*, pp. 221–273, 2001.

[22] A. M. Bloch, *Nonholonomic Mechanics and Control*. Springer New York, NY, 2003.

[23] S. D. Kelly and R. M. Murray, "The geometry and control of dissipative systems," in *Proceedings of 35th IEEE Conference on Decision and Control*, vol. 1. IEEE, 1996, pp. 981–986.

[24] S. Revzen and J. M. Guckenheimer, "Estimating the phase of synchronized oscillators," *Physical Review E*, vol. 78, p. 051907, Nov 2008.

[25] S. Revzen and M. Kvalheim, "Data driven models of legged locomotion," in *Micro- and Nanotechnology Sensors, Systems, and Applications VII*, T. George, A. K. Dutta, and M. S. Islam, Eds., vol. 9467, International Society for Optics and Photonics. SPIE, 2015, p. 94671V.

[26] E. M. Purcell, "Life at low reynolds number," *American journal of physics*, vol. 45, no. 1, pp. 3–11, 1977.

[27] "Abaqus/standard, version 2021," United States, 2021.

[28] S. A. Chester and L. Anand, "A thermo-mechanically coupled theory for fluid permeation in elastomeric materials: Application to thermally responsive gels," *Journal of the Mechanics and Physics of Solids*, vol. 59, no. 10, pp. 1978–2006, 2011.

[29] Y. Ding, S. S. Sharpe, K. Wiesenfeld, and D. I. Goldman, "Emergence of the advancing neuromechanical phase in a resistive force dominated medium," *Proceedings of the National Academy of Sciences*, vol. 110, no. 25, pp. 10 123–10 128, 2013.

[30] M. S. Madhav and N. J. Cowan, "The synergy between neuroscience and control theory: the nervous system as inspiration for hard control challenges," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 243–267, 2020.

[31] A. Hasan, J. Saliba, H. Pezeshgi Modarres, A. Bakhaty, A. Nasajpour, M. R. Mofrad, and A. Sanati-Nezhad, "Micro and nanotechnologies in heart valve tissue engineering," *Biomaterials*, vol. 103, pp. 278–292, 2016.

[32] M. S. Erden and K. Leblebicioğlu, "Free gait generation with reinforcement learning for a six-legged robot," *Robotics and Autonomous Systems*, vol. 56, no. 3, pp. 199–212, 2008.

[33] V. Tsounis, M. Alge, J. Lee, F. Farshidian, and M. Hutter, "Deepgait: Planning and control of quadrupedal gaits using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3699–3706, 2020.

[34] A. W. Winkler, C. D. Bellicoso, M. Hutter, and J. Buchli, "Gait and trajectory optimization for legged systems through phase-based end-effector parameterization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1560–1567, 2018.