

Learning physics-informed simulation models for soft robotic manipulation: A case study with dielectric elastomer actuators

Manu Lahariya^{1*}, Craig Innes², Chris Develder¹ and Subramanian Ramamoorthy²

Abstract—Soft actuators offer a safe, adaptable approach to tasks like gentle grasping and dexterous manipulation. Creating accurate models to control such systems however is challenging due to the complex physics of deformable materials. Accurate Finite Element Method (FEM) models incur prohibitive computational complexity for closed-loop use. Using a differentiable simulator is an attractive alternative, but their applicability to soft actuators and deformable materials remains under-explored. This paper presents a framework that combines the advantages of both. We learn a differentiable model consisting of a material properties neural network and an analytical dynamics model of the remainder of the manipulation task. This physics-informed model is trained using data generated from FEM, and can be used for closed-loop control and inference. We evaluate our framework on a dielectric elastomer actuator (DEA) coin-pulling task. We simulate the task of using DEA to pull a coin along a surface with frictional contact, using FEM, and evaluate the physics-informed model for simulation, control, and inference. Our model attains $\leq 5\%$ simulation error compared to FEM, and we use it as the basis for an MPC controller that requires fewer iterations to converge than model-free actor-critic, PD, and heuristic policies.

Index Terms—Dielectric elastomer actuators, Differentiable simulator, Finite element methods, Model predictive control, Neural Networks, Physics based machine learning, Soft Actor-Critic, Soft robotics

I. INTRODUCTION

Soft robotic actuators provide a safe, adaptive, low-cost solution for movement tasks such as grasping and dexterous manipulation [1]. Precision manipulation using soft actuators however is a major challenge, as it requires modelling the deformable actuators’ dynamic within the context of the manipulation task [2]. Such models are then used to learn accurate control strategies via simulation [3]. Recently *differentiable simulators* have been used to learn controllers in closed-loop scenarios by allowing the use of gradient-based optimization methods (e.g., Model Predictive Control, MPC) [4]. They have also been used for inference and data generation tasks.

Simulating deformable robots and contact rich manipulation is expensive [3]. Traditional methods model such dynamics by decomposing their geometry. For example, *Position Based Dynamics* approximates multi-body physics by deconstructing the system into particles [5]. However, these methods fail to accurately capture the true underlying physics, making it difficult to meaningfully interpret or




Fig. 1. Coin pulled by a dielectric elastomer actuator (DEA)—a soft actuator that deforms under electric actuation. Our framework learns accurate physics-informed differentiable simulators and model-based control for such soft robot manipulation.

constrain the particles. The continuum mechanics and contact dynamics of deformable materials are difficult to model with such approximate methods, leading to physically unrealistic results, which hinders model-based control. Physically accurate simulation of soft robotic manipulation requires modelling the underlying equations, defined by complex Ordinary/Partial Differential Equations (ODEs/PDEs).

Finite Element Methods (FEMs) provide a numerical method for solving such equations. Yet despite the ability of FEMs to accurately model such phenomena, integrating FEM simulation with closed-loop control is challenging due to their computationally expensive meshing: unless the meshes are dense and cover the domain, fidelity is poor.

This paper’s key idea is to generate data from an accurate (but slow) FEM model to learn an approximate (but fast) physics-informed model f for soft robotic manipulation. Our framework uses f as a differentiable simulator for simultaneous closed-loop control and inference. Our model f is composed of two parts: a *material network* m — which approximates deformable material behaviour (e.g., hyperelastic) — and *dynamics* d — equations representing the physical context of manipulation task (sliding motion under frictional contact with the surface).

Our framework’s objective is to learn fast physics-informed models that can be evaluated in real-time without significant loss in accuracy, and to use these models for control. As an illustrative use case, we focus on soft robotic pulling task using Dielectric Elastomer Actuators (DEAs). DEAs are soft actuators made using electroactive polymers that convert electrical work to mechanical work via expanding or bending motion. In our task, the goal is to pull a stationary coin by deforming the free end of the DEA (Fig. 1). We learn the physics-informed model for this pulling motion f , and evaluate its accuracy as a simulator against FEM simulations. For control, we use the differentiability of f to learn a model-based control policy (with the solver

¹Authors are with IDLab, Ghent University – imec, Technologiepark-Zwijnaarde 126, 9052 Ghent, Belgium,

²Authors are with the School of Informatics, University of Edinburgh, 10 Crichton St, EH8 9AB, United Kingdom,

*Corresponding Author: e-mail: manu.lahariya@ugent.be.

from [6] and infer the parameters of the system’s dynamics.¹

Our main contributions are: (i) a closed-loop control framework for soft robotic manipulation, that uses a differentiable physics-informed model f trained using FEM (Section II), (ii) the design of an exemplary DEA pulling task (Section III), that is simulated in FEM (Section IV), and (iii) performance evaluation of model f and its use in closed-loop control. For the latter, we compare simulation accuracy of f with both a FEM model and a baseline neural network. Additionally, we compare the model-based control policy (MPC with f), with (i) a model-free control policy (learnt using the Soft Actor-Critic SAC algorithm [7]), (ii) a PD control policy (evaluated previously for DEA control [8]), and (iii) a heuristic control policy² (inspired by typical soft-robotic control policies [3]). We design experiments (Section VI) across 8 DEA pulling setups to evaluate our framework and answer the following questions:

- (Q1)** How to define f using the physical laws of the system? What is the simulation accuracy of f in a system with new *unknown* parameters (e.g., frictional coefficient)?
- (Q2)** What is the performance of model-based control policy (based on f), compared to other control policies?
- (Q3)** What is the accuracy of the inferred model parameters?

Our results show f provides $\leq 5\%$ simulation error compared to FEM. Further, in closed-loop control, an MPC using f outperforms all other policies, while simultaneously inferring system properties with $\leq 10\%$ error. Videos of the manipulation task and other supplementary materials are available at <https://sites.google.com/view/phy-informed-sim-soft-robot/home>.

A. Related Work

Soft robots are inspired by biological systems, e.g., where animals use muscles to achieve safe actuation and control [1]. Engineers use soft actuators to develop similarly safe, quick, adaptable, and precise robotic manipulation [2]. These soft actuators generate mechanical work under a specific actuation, e.g., shape memory alloys respond to thermal actuation, hydraulic actuators respond to pressure, etc. Learning control for soft actuators requires accurate simulation models that are used inside the control loop [4]. Designing simulator models for soft actuators is a challenging task, traditionally using particle based models (e.g., liquids [9]). In recent years, researchers are using FEM modelling that allows highly accurate modelling of deformable materials (e.g., fabric [10], composite materials [11]).

Dielectric elastomers (DE) are electroactive polymers that produce deformation under the influence of an external electric field. DEA are soft actuators that use thin layers of DE materials to achieve actuation under the stimulus of electric activation. DEAs provide fast and large deformation, are lightweight, and have a high energy density, which makes them promising candidates for soft robotic applications [12].

¹For the case study of DEA pulling, coin mass m_c and kinetic friction coefficient μ_b are inferred. For details, refer to Section III-B.

²The heuristic policy linearly ramps up actuation voltage until the ‘episodic’ task terminates. For details, refer to Section VI

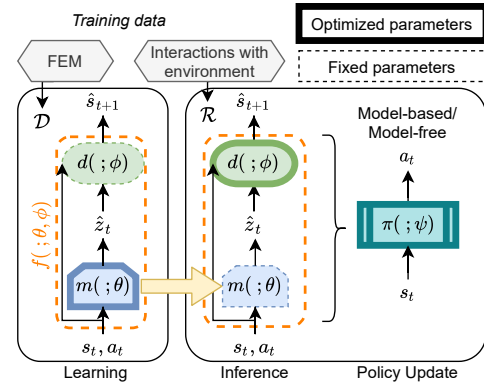


Fig. 2. Training model f - the simulator for our task. Material network m is optimized during learning; dynamics d inferred via closed-loop interaction.

Hence, DEA has been explored to design soft robotic grippers [13], underwater robots [14], crawling robots [15], etc.

There have been several previous approaches to modelling DE behaviour. A mathematical fractional Kelvin-Voigt model for DEA is presented in [8], where they eliminate overshoot in PID control. In [16], a spring-dashpots model describes the DEA’s dynamics, with model parameters determined from the response of the robot. FEM models for DE material have been explored, where deformation in unimorph DEA with inhomogeneous geometry modeled in [17] uses piezoelectric elements. A simplified finite element analysis of a dielectric bending actuator is performed in [13]. The FEM model for a gripping actuator in [18] is based on a custom defined material. These methods focus on modelling the DEA behaviour in isolation (separate from the manipulation task). Thus, they lack an understanding of the task context in which the DEA manipulator is being used. In contrast, our method simulates the complete manipulation task, allowing us to learn both DEA behaviour and the task specific context.

An accurate simulator of the manipulation task can assist in learning a controller. For example, a position based dynamics simulator (defined using particle interactions) is used in [9] to develop control strategies for pouring liquid. A FEM based differentiable simulator in [19] is used to learn control strategies for cutting. The above methods are designed to control one specific manipulation task (e.g., in [20], the model is explicitly engineered for cutting). In contrast, our control framework can be used for *any* robotic manipulation task that can be simulated in FEM.

Learning based modelling and control can be used for soft robot manipulation. For example, a differentiable model of a soft robot’s quasi-static physics is learned in [21], and then used to perform gradient-based optimization to perform open-loop control. In [22], a recurrent neural network is used as a fixed forward model within a policy learning algorithm for closed-loop soft robot control. In contrast, our trained physics-informed model f can infer properties of new *unknown* setups, adapting during the closed-loop control process.

II. CONTROL FRAMEWORK

The objective of our control framework is to learn a control policy π for a manipulation task. Figure 2 shows the closed-loop control design using the trained physics-informed model f of the manipulation task along with policy π .

A model-based control approach utilizes a forward model of the system: $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, where \mathcal{S} is the state space, and \mathcal{A} is the control action space. For each timestep t , the state is $s_t \in \mathcal{S}$, and the control action is $a_t \in \mathcal{A}$. For MPC, the optimal control actions at each timestep is estimated by solving the optimization problem defined in Eq. (1), where a_{init} is the initial action. We particularly choose to use a physics-informed f , which is differentiable, and allows us to use gradient-based methods to solve this optimization problem (such as finite-horizon Linear Quadratic Regulator, LQR [6]). The objective of the control (e.g., get to a target location) is used to define the cost function $\mathcal{C} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ (e.g., distance from the target location). For example, in DEA pulling, cost function \mathcal{C} is defined for the objective of achieving target state (i.e., the target location of the coin) with penalty on control actions to minimize actuation voltage of the DEA. Eq. (2) shows the cost function, where s_t^* is target state, and w_s, w_a the state and action penalties.

$$\arg \min_{s_{1:T} \in \mathcal{S}, a_{1:T} \in \mathcal{A}} \sum_{t=1}^T \mathcal{C}(s_t, a_t) \quad (1)$$

$$\text{s.t. } s_{t+1} = f(s_t, a_t) \quad \text{and} \quad a_1 = a_{\text{init}}$$

$$\mathcal{C}(s_t, a_t) = \frac{1}{2}(w_s(s_t - s_t^*)^2 + w_a a_t^2) \quad (2)$$

The physics of the robotic manipulation task includes, (i) the physical laws of the deformable material behaviour, e.g., electromechanical/hyperelastic behaviour, characterized by high order ODEs/PDEs that are computationally complex, and, (ii) the physical laws built according to the context of the manipulation task, e.g., sliding motion laws, or gravity. In modeling these physics, interaction variables z are introduced to describe the contact properties (e.g., force, stress, pressure) between the deformable material and its surroundings. In particular, for DEA pulling, z are the forces exerted by the DEA actuator on the contact surface.

We simulate the manipulation task using a FEM model to numerically solve the associated physics equations. In this model, state s_{t+1} and interaction variables z_t are simulated, given s_t and a_t . The FEM model for DEA pulling is described in Section IV. How the simulated data is then used to train a physics-informed model f is described below.

A. Physics-informed model (f)

The physics-informed model f has two parts:

- (i) The material network (m): A function approximator with weights θ estimating interaction variables \hat{z} (Eq. (5)). These interaction variables (z) characterize deformable material behaviour in the manipulation task (e.g., forces by DEA on contact surface).
- (ii) The dynamics (d): Physical laws characterizing the motion/dynamics of the system in the form of mathematical

Algorithm 1 Learning

Output: Material network m ;

- 1: Randomly initialize weights θ , and fix parameters ϕ ;
 - 2: **while** not stopping condition **do**
 - 3: $a_t \leftarrow$ select action at t using a fixed policy;
 - 4: $s_{t+1}, z_t = FEM(s_t, a_t)$;
 - 5: Dataset $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, a_t, z_t, s_{t+1})$;
 - 6: **end while**
 - 7: Using all data from $\sim \mathcal{D}$; // Say $|\mathcal{D}| = N$ data samples
 - 8: $\hat{z}_t = m(s_t, a_t; \theta)$;
 - 9: $\hat{s}_{t+1} = f(s_t, a_t; \theta, \phi)$;
 - 10: $\theta \leftarrow \theta - \alpha \frac{d\mathcal{L}_l(\theta, \phi)}{d\theta}$; // Update θ using \hat{z} and \hat{s} , Eq. (6)
 - 11: **return** m
-

equations (e.g., linear equations or ODEs/PDEs representing sliding or gripping). The dynamics d estimate the next state using interaction variables z , state s , and action a (Eq. (4)). Parameters ϕ describe the system's physical properties, e.g., the mass of coin.

Thus, we can write the model f as in Eq. (3).

$$\hat{s}_{t+1} = f(s_t, a_t; \theta, \phi) \quad (3)$$

$$f(s_t, a_t; \theta, \phi) = d(\hat{z}_t, s_t, a_t; \phi) \quad (4)$$

$$\hat{z}_t = m(s_t, a_t; \theta) \quad (5)$$

where θ are material neural network weights, and ϕ the parameters (e.g., mass) for system dynamics. The material model m of a deformable material can describe an actuator (e.g., DEA), or the manipulated object (e.g., cloth) depending on the manipulation task. The interaction variables z depend only on the deformable material (i.e., the material network m), and are not impacted by the dynamics d of the manipulation. For example, in DEA pulling, m describes the actuator behaviour of a unimorph DEA (Section IV). Our model f is a physics informed neural network [23], where physical rules can be imposed on interaction variables z in the form of ODEs/PDEs. However, for the case study of DEA pulling, a system of linear equations sufficiently defines dynamics d .

We next show how to define f using physical laws. Model f is differentiable, and captures the manipulation task physics through dynamics d . In addition to its use as a simulator to generate data, we also use it for inference, and for gradient-based control learning.

B. Training and policy synthesis

The physics-informed model f and policy π are learnt in two steps. First, a *Learning* step optimizes weights θ of m , using data generated by the FEM model of the task. Second, a *Control* step, where policy π is learnt via environment interactions. These interactions are used to infer parameters ϕ (e.g., coin mass) of dynamics d , informing f . We then use f to learn π .

a) Learning

The weights θ are optimized by minimizing the loss function based on the error in estimating material model,

Algorithm 2 Control

Input: Trained weights θ

- 1: Randomly initialize parameters ϕ , and ψ , and fix θ ;
 - 2: $s_1 \leftarrow env.reset()$;
 - 3: **while** not stopping condition **do**
 - 4: $a_t = \pi(s_t; \psi)$;
 - 5: $s_{t+1}, r_t \leftarrow env.step(a_t)$;
 - 6: Replay buffer $\mathcal{R} \leftarrow \mathcal{R} \cup (s_t, a_t, r_t, s_{t+1})$;
 - 7: **if** it's time to update **then**
 - 8: Randomly sample B transitions from $\sim \mathcal{R}$;
 - 9: // Inference
 - 10: $\hat{s}_{t+1} = f(s_t, a_t; \theta, \phi)$;
 - 11: $\phi \leftarrow \phi - \alpha_\phi \frac{d\mathcal{L}_i(\theta, \phi)}{d\theta}$; // Update ϕ using \hat{s} , Eq. (7)
 - 12: // Policy Update
 - 13: Update ψ by policy defined updates, e.g., SAC [7];
 - 14: **end if**
 - 15: **end while**
-

i.e., m : $(z_t - \hat{z}_t)$, and the error in enforcing dynamics, i.e., d : $(s_t - \hat{s}_t)$. Incorporating the loss encountered in \hat{s} ensures that our model adheres to dynamics d (as shown by optimization of physics informed neural networks [23]). Algorithm 1 shows how θ is optimized by fixing parameters ϕ and minimizing the learning loss \mathcal{L}_l (Eq. (6)). A fixed policy is used to select actions a_t (e.g., a random or uniform policy). The learning rate for weights θ is α_θ and number of data samples is N .

$$\mathcal{L}_l(\theta, \phi) = \frac{1}{N} \sum_{t=1}^N (z_t - \hat{z}_t)^2 + \frac{1}{N} \sum_{t=1}^N (s_t - \hat{s}_t)^2 \quad (6)$$

b) Control

Algorithm 2 shows our steps for closed-loop control. First, in inference, parameters ϕ are optimized by minimizing the dynamics estimation error, i.e., d : $(s_t - \hat{s}_t)$. Eq. (7) gives the loss function, with learning rate γ_ϕ and batch size B .

$$\mathcal{L}_i(\theta, \phi) = \frac{1}{B} \sum_{t=1}^B (s_t - \hat{s}_t)^2 \quad (7)$$

Second, we learn policy π with weights ψ . Updates in ψ are defined using the underlying policy π and its training objective. For the model-free policy trained using the SAC algorithm, ψ is updated based on the loss function from [7]; the weights ψ represent weights of the neural network. For model-based MPC, we use our trained model f and gradient-based optimization to estimate the best action. Thus, with MPC, updating parameters ψ (line 13) is unnecessary.

III. SOFT ROBOTIC DEA PULLING

We design the manipulation task of coin pulling using a unimorph Dielectric Elastomers Actuators (DEAs) to evaluate the framework proposed in Section II. The deformable DEA actuator is made of Dielectric Elastomers (DEs), which are a type of electroactive polymers that produce mechanical strain under the influence of electric voltage. Thus, a DE




Fig. 3. (a) Unimorph DEA: Active layer expands under influence of electric voltage, causing bending. (b) Problem setup: DEA moves stationary coin.

TABLE I
STATE AND ACTION DEFINITION FOR TIMESTEP t

s_t	x_t	Location of coin along x -axis at time t
	u_t	Velocity of coin along x -axis at time t
a_t	V_t	Voltage applied on the DEA
	Δt	time difference between t and $t+1$
z_t	$F_{x,t}$	Force along x -axis by DEA on coin c
	$F_{y,t}$	Force along y -axis by DEA on coin c

membrane expands its area when a voltage is applied across its thickness [24].

Figure 3(a) shows a unimorph DEA, with one active and one constraining layer. The active layer expands under externally applied voltage causing the bending motion. The DEA is fixed at one end, and the other end rests freely on a circular coin c . On actuation, the DEA acts as a soft robotic finger, pulling the coin. A controller policy π can be learnt to achieve a certain displacement in the coin. Figure 3(b) shows the 2D view of the setup, where the mass of the coin is m_c , the kinetic friction coefficient between the coin and DEA is μ_t , and the kinetic friction coefficient between the coin and bottom surface is μ_b . The displacement of the coin depends on such parameters of the system. A pulling coin setup C_c is characterized by fixed values of $\{m_c, \mu_t, \mu_b\}$. Setups C_1, C_2, \dots represent pulling different coins, based on different parameter values.³

The physics-informed model f of the system is defined by the variables shown in Table I. The state of the system at time-step t is characterized by the location x_t and velocity u_t of the coin along the x -axis. The action comprises the voltage (V_t) applied on the DEA and Δt ,⁴ and the hidden variables are the forces (F_x and F_y) applied by the DEA on the top surface of the coin.

A. Material network (m)

Modeling non-linear properties of DEs require modeling the effects of hyperelasticity and Maxwell stress [24]. On application of voltage V , maxwell stress causes the bending actuation in DEA. The actuated DEA exerts forces F_x and F_y on the top surface of the coin, which results in its motion. The material network used to estimate these forces is defined in Eq. (8). We simulate DEA pulling using FEM, to generate

³We consider the coins to be of fixed dimensions (i.e., fixed volume), and thus change the mass m_c by changing the density ρ of the coin material. For further details, please refer to Section VI-A.

⁴Note that in FEM, the time between successive simulation datapoints may vary. Thus, Δt is variable during training the model f .




Fig. 4. FEM of DEA pulling, (a) inactive DEA with dimensions, (b) active DEA during static friction stage (c) motion occurs in kinetic friction stage.

data and optimize weights θ (Section IV).

$$\hat{F}_{x,t}, \hat{F}_{y,t} = m(x_t, u_t, V_t, \Delta t; \theta) \quad (8)$$

B. Dynamics (d)

Physical laws of the pulling setup define the system dynamics d (Section II-A). There are two stages during pulling: *static friction* (forces are applied but there is no motion), and *kinetic friction* (applied forces cause motion in coin). An actuation threshold voltage V^T is required to achieve a minimum coin displacement (i.e., to get to the stage of kinetic friction). Coin acceleration A_t is due to the net force in the x -axis (Eq. (9)), where F_μ is the opposing frictional force. We calculate F_μ using Eq. (10), assuming linear growth in frictional force during the stage of static friction, and a no-slip condition on the top surface.

$$\hat{F}_{x,t}^{Net} = m_c A_t = \hat{F}_{x,t} - F_{\mu,t} \quad (9)$$

$$F_{\mu,t} = \begin{cases} \mu_b (\hat{F}_{y,t} + m_c g) & \text{if } V \geq V^T; \\ \frac{\mu_b V_t}{V^T} (\hat{F}_{y,t} + m_c g) & \text{otherwise} \end{cases} \quad (10)$$

where g is gravitational acceleration (9.8 m/s^2). We assume a frictional velocity decay for a moving coin if DEA actuation is stopped (i.e., $V_t = 0$). These dynamics disregard non-linear motion in coin (with high DEA actuation voltages, where DEA loses contact with the coin surface). The dynamics d of this pulling setup is characterized by parameters ϕ (Section II), which are: coin mass m_c and frictional coefficient μ_b . While training m , these values are fixed. We infer ϕ during closed-loop environment interactions.

The next location and velocity of the coin, i.e., \hat{x}_{t+1} and \hat{u}_{t+1} , are calculated using A_t , x_t , u_t , and Δt and equations of motion. Thus, the dynamics d is a set of linear equations based on the laws of motion.

IV. FEM OF DEA PULLING

FEM is a numerical method for solving differential equations of a physical system. Its complexity depends on both

the task and available computational power. The physical system of DEA pulling consists of a unimorph DEA on the fixed surface and a solid coin (Section III). We use commercially available software ABAQUS [25] to build a 3D model of DEA pulling. Figure 4 shows the simulated setup, during the starting stage and the stages of static and kinetic friction. In the stage of static friction, we clearly see no motion in the coin even when actuating the DEA.

a) DE Material

To model the non-linear elastic behaviour in DE membranes, recent works use hyperelastic material models, like the Gent and the Neo-Hookean models [26]. For the small strains present in the thin membrane of our DEA ($\leq 10\%$), such hyperelastic behaviour reduces to linear elasticity. Furthermore, since no commercial FEM packages provide DE elements out of the box, we approximate the behaviour of our DE material in FEM using piezoelectric materials elements [27]. We modify the piezoelectric finite elements material properties to model the Maxwell stress effect observed in dielectric materials. The Maxwell stress p on the DE membrane is given by Eq. (11) [24]. Similar to DE, piezoelectric materials exhibit strain when in the presence of electric fields. The piezoelectric stress is given by Eq. (12).

$$p = e_0 e_r \left(\frac{V}{z} \right)^2 \quad (11)$$

$$\sigma_{ij} = D_{ijkl}^E \epsilon_{kl} - e_{mij} E_m \quad (12)$$

where V is the applied voltage across thickness z of the DE membrane, the relative permittivity is e_r , and the permittivity of free space is e_0 . The piezoelectric elastic stiffness matrix is D_{ijkl}^E , the strain tensor is ϵ_{kl} , the stress coefficient is e_{mij} and the electric potential gradient is E_m . As detailed in [17] the piezoelectric stress becomes approximately equal to the Maxwell stress for a thin membrane, such that the strain e_{zz} in the direction of thickness (z -axis) is given by:

$$e_{zz} = e_r e_0 E_z \quad (13)$$

A. FEM simulation settings

Figure 4(a) shows our FEM setup when the DEA is inactive (zero voltage). The mesh consists of an 8-node linear brick (ABAQUS element type C3D8E). Each DE membrane has 10 elements.⁵ For meshing the coin, ABAQUS's internal meshing strategy is used to generate 20 elements.

For both the active and constraining layers of the DE material, the Poisson's ratio is 0.5 and Young's modulus is 0.56 MPa [28]. We use Eq. (13) to calculate $e_{zz} = 3.68$. All other piezoelectric coefficients are zero. We assume elastic behaviour for the coin. The bottom surface and the fixed end of the DEA are constrained using encastre boundary condition. The top surface assumes a no-slip condition.

The coin rests on the frictional surface with coefficient $\mu_b \in \{0.2, 0.25\}$ (defined as tangential behaviour in the contact interactions in ABAQUS). Similarly, the free end of the DEA rests atop the coin with $\mu_t \in \{0.5, 0.55\}$. To

⁵We limit the number of mesh elements due to software limitations.

simulate a real scenario, we include gravitational load ($g = 9.8 \text{ m/s}^2$). We also do not assume a no-slip condition between the top of the coin and DEA, in contrast to the dynamics in Section III-B. Coin mass (m_c) is calculated using its volume and density $\rho \in \{7.7, 7.8\} \text{ g/cm}^3$, thus, making $m_c \in \{1.36, 1.38\} \text{ g}$. The total time simulated in FEM is 1 s, with Δt between points determined by the internal solver.

V. PARAMETERS SETTINGS

This section details the architectures, parameters, and training procedures used to evaluate our framework. An experimental setup C is defined using differing values of $\langle \mu_t, \mu_b, m_c \rangle$ (Section III). We use 8 setups, denoted $\{C_1, C_2, \dots, C_8\}$. An FEM model is developed for each setup. To collect the dataset for each model, we apply a linearly increasing electric potential load ($V_t \in \{0.0, 400.0\}V$) to the top surface of the active DE layer. Table I describes the values collected at each timestep t . Each dataset contains 1000-2000 data points. Across setups, $\langle x_0, u_0 \rangle = \langle 0.0, 0.0 \rangle$.

During learning, we initialize parameters ϕ (m_c and μ_b , Section III-B) of dynamics d using the true values from the FEM model. Note that μ_t is *not* used as a parameter in the dynamics of f , but, is required for the FEM modeling. For each setup, the threshold voltage (V^T) is the voltage required to achieve a displacement of -10^{-5} m . We assume the coin loses contact with the DEA for $V_t \geq 300 \text{ V}$.

Physics-informed model f is developed using Pytorch [29] and contains the material network m and the dynamics d (Section II). The material network m is a fully connected neural network with four input nodes, two output nodes, three hidden layers with 128 neurons each, and rectified linear (ReLU) activation functions. An ADAM optimizer with learning rate 0.001 is used to minimize Eq. (6) and Eq. (7) for 1,000 iterations. An early stopping criterion is used based on validation loss with 0.0 minimum change. The baseline neural network uses ReLU activations, while the recurrent neural network uses *LSTM* cells. The *LSTM* uses an 8-step recurrence.

The target state for the controller is $x^T = -1.0 \text{ mm}$, i.e., goal is to achieve a 1 mm displacement. The manipulation ‘episode’ is simulated, and terminates when $\|x_t - x^T\| < 0.01 \text{ mm}$. To simulate real-world discrepancies, we add gaussian noise to simulated state x_t with zero mean and standard deviation of 0.001 m. We average results for 10 ‘episodes’ for all controllers. Batch size (Algorithm 2) is 256. During control, the time increment Δt is fixed to 0.001 s—the policy does not provide it as an output. For inference, m_c is initialized to 0.001 g and μ_b is initialized to 0.2.

We use the model-based control policy from [6]. For MPC, we set the number of timesteps 20, LQR iterations to 20, and action penalty to 0.001. The model-free policy is trained using Soft Actor-Critic (SAC) [7]. For SAC, fully connected neural networks are used for actor and critic with two hidden layers of 256 neurons each. The value of τ (soft updates) is set to 0.005, and networks are optimized using MSE loss and the ADAM optimizer. For the PD controller, the value of K_p is set to -0.5 and K_d is set to 5.

TABLE II
MEAN ABSOLUTE ERRORS IN SIMULATION

Simulated value	f	LSTM	NN
x_t ($1 \times 10^{-4} \text{ m}$)	0.34	3.86	9.86
$F_{x,t}$ ($1 \times 10^{-4} \text{ N}$)	3.68	—	—
$F_{y,t}$ ($1 \times 10^{-4} \text{ N}$)	12.01	—	—

VI. EXPERIMENTAL RESULTS

This section presents experiments describing the high simulation and inference accuracy of f , while also developing an effective closed-loop soft robotic controller. For the case of DEA pulling, we evaluate, (i) the accuracy of the f as a simulator, and (ii) the accuracy of f in inference, and (iii) the closed-loop MPC controller that utilizes f .

A. Simulation (Q1)

We design an experiment to answer Q1, i.e., to show that f can simulate data for new parameter settings, we train and simulate on different setups (defined in Section III). Note that different setups represent different coins, with different mass m_c and frictional coefficients μ_b and μ_t .

A simulation experiment set has coin setups given by $\{C_{\text{train}}, C_{\text{val}}, C_{\text{test}}\}$. FEM data from $\{C_{\text{train}}, C_{\text{val}}\}$ is used to optimize weights θ (material network). The objective of learning step (Algorithm 1) is to optimize weights θ of the material network m . During training and validation, the material network m is same across coin setups (i.e., same θ for C_{train} and C_{val}) and dynamics d are specific to coin setups (i.e., ϕ based on C_{train} and C_{val}). During testing, data is simulated recursively (for $T = 1 \text{ s}$) in a test setup C_{test} , using simulator f . This f consists of (i) a material network with previously optimized θ , and (ii) the dynamics with parameters ϕ_{test} (based on C_{test}). We present averaged results for 6 experiment sets (E.g., for first set, $\{C_{\text{train}}, C_{\text{val}}, C_{\text{test}}\}$ are $\{C_1, C_2, C_3\}$, the second are $\{C_2, C_3, C_4\}$, etc.).

We evaluate the absolute errors encountered in simulating \hat{s}_{t+1} (x_t and u_t) at each timestep t . For example, error in x_t is given by $e_t^x = |\hat{x}_t - x_t|$, where \hat{x}_t is the location simulated using f , and x_t is the true value (from the FEM dataset). Absolute errors in data simulated using a black-box baseline Neural Network (NN) and long short term memory recurrent neural network (*LSTM*) trained using data from C_{train} and C_{val} are also included (e.g., a NN that simply approximates $\hat{s}_{t+1} = \text{NN}(s_t, a_t; w)$).

Table II provides the mean absolute errors in simulating location (x_t) and forces using physics-informed model f and the baseline models across all test setups. Figure 5 shows the absolute error in x for all test setups for f , NN, and *LSTM*. In all cases, f outperforms the baseline models. Additionally, the average absolute error in x simulated using f is less than 0.05 times the magnitude of the actual values, i.e., we note approximately $\leq 5\%$ error compared the FEM simulation.

In the first half of the simulation time during static friction (Section III-B), we see negligible displacements and increasing F_y . In the latter half, the kinetic frictional force becomes stable, and we see a change in coin location (x_t). Our model




Fig. 5. Absolute error in simulated location x_t . Each box point has data for all simulation points.

f accurately estimates x in both stages. We see a similar accuracy for f compared to the FEM, NN , and $LSTM$ in simulating velocity u_x for the coin. While decreasing the neurons per layer in f from 128 to 64 (Section V) results in no statistically significant increase in average error, decreasing the capacity of baseline NN by the same amount results in an increase from 15% average error to 50%.

Real-time control in manipulation tasks requires fast simulations. For 1 second of simulation with 700 points, the FEM model takes 130 seconds (average for all test coins) — a prohibitive duration for real-time control. In comparison, our physics-informed model f takes ≤ 0.7 seconds. While both baseline models are also faster than FEM, (0.3 and 1.2 seconds for NN and $LSTM$ respectively), they provide poor accuracy compared to f . Our proposed model f is fast as opposed to FEM but provides high fidelity to FEM as a simulation model of soft robotic manipulation.

The absolute error in forces F_x and F_y in the region of static friction is higher compared to the region of kinetic friction. This happens because we optimize material network m using a physics informed loss function, i.e., a loss function that is based on the error in the next state s and the error in the interaction variables z (Eq. (6)). Optimizing m using this loss function assists in learning the overall manipulation task behaviour, as opposed to only learning the outputs of m (F_x and F_y). This behaviour is non-restrictive, as the objective of our model is to learn the next state of the motion, which is simulated accurately.

B. Control (Q2)

In this experiment, we evaluate f to answer Q2 (performance of model-based MPC compared to other control policies?). In the control step (Algorithm 2), we learn closed-loop control for test setups C_1 and C_2 . Prior to this, in the learning step (Algorithm 1), we train the model f_a using the FEM data from setups $\{C_3, \dots, C_8\}$. We do not use the data from C_1 and C_2 during learning to avoid information leakage. For C_1 and C_2 , we learn the following policies:

- (i) *MPC policy*: A model-based control policy defined using differentiable model f and an MPC solver [6]. The MPC uses the model with inferred physical parameters.
- (ii) *SAC policy*: A model-free Actor-Critic policy learnt using the Soft Actor-Critic algorithm [7],
- (iii) *PD policy*: A feedback based control policy (previously tested for DEA control [8]),




Fig. 6. DEA pulling control for C_2 . Terminal state: coin is ≤ 0.01 mm from x^T (Solid line: average for 10 ‘episodes’, shaded region: 25-75%)

- (iv) *Heur policy*: A heuristic control policy that linearly ramps up actuation voltage (i.e., voltage increases by 0.5 V after each iteration until terminal state).

The *Heur policy* is inspired by typical soft-robotic control policies [3]. The environment is simulated by trained physics-informed models $f_{s,1}$ and $f_{s,2}$. This is due to the lack of a real-world DEA setup (However, Section VI-A shows our physics-informed models are accurate simulators).

Figure 6 shows the average coin location x_t during closed-loop control of test setup C_2 . Average is calculated across 10 episodes to compare the *MPC policy* (model-based), with *SAC policy* (model-free), a *PD policy* and a *Heur policy*. We notice similar results for both test setups C_1 and C_2 . The variance in coin position is due to the noise added in the simulated observations to estimate a realistic scenario.

The coin reaches the target in ≤ 200 iterations under the *MPC policy*, i.e., ≤ 200 observations are captured in the control loop. In contrast, *SAC policy*, *PD policy*, and *Heur* capture approximately 10,000, 1500, and 500 iterations. Thus, during control, the *MPC policy* reaches the target in the least number of steps, where the computation time per iteration (average 1.2 s) is not computationally prohibitive for a real world manipulation. We note similar computation times for the model-free *SAC policy* (average 0.005 s) and *PD policy* (average 0.001 s).

During control using *Heur policy* we notice sudden motion towards x^T after approximately 280 iterations. This represents the transition from stage of static friction to stage of kinetic friction. The *Heur policy* linearly ramps up actuation voltage every iteration, and thus, does not depend on location feedback. In contrast, both *PD policy* and *SAC policy* rely on observed location, and take longer to reach and manipulate the coin in the stage of kinetic friction.

C. Inference (Q3)

This experiment evaluates how accurately f infers the model parameters. Algorithm 2 infers m_c and μ_b for test setups C_1 and C_2 . Similar to the control experiment, in the learning step, we train f_a using FEM data from setups $\{C_3, \dots, C_8\}$. The mass m_c and frictional coefficient μ_b inferred by physics-informed model f are compared against their real values (used in C_{test} during control with *MPC policy*). Inferred m_c converges to $\leq 10\%$ error compared to the real value within 2,000 iterations, while μ_b , the converges within 300 iterations. Similar results hold across both test setups and all episodes.

VII. CONCLUSIONS

This paper presents a framework to learn a differentiable simulator to develop a corresponding controller for soft robotic manipulation. We defined a physics-informed model f consisting of a material network m , and dynamics d . This model f can be used as a simulator for data-generation, inference, and control policy optimization. We designed a soft-robotics case study where a coin is pulled using unimorph DEA. FEM simulation of the DEA generated data to train f . Our experiments used multiple setups to evaluate the framework in learning f and model-based control.

From our analyses, we conclude, (i) the physics-informed model f trained using the proposed framework can simulate new setups (characterized by parameters ϕ) with $\leq 5\%$ error compared to FEM (Fig. 5); (ii) a closed-loop MPC policy based on differentiable model f outperformed all other policies in orders of hundreds of iterations (Section VI-B); (iii) f can be used for accurate inference of the parameters ϕ : m_c and μ_b (Section VI-C). Future research directions include: Further evaluating this control framework with *physical* soft robotic actuators; Evaluating the model f in tasks with large degrees of freedom; and exploring model-based policies with lower computational requirements compared to MPC.

ACKNOWLEDGMENT

S. Ramamoorthy would like to acknowledge financial support from the Alan Turing Institute, for the project ‘Enabling advanced autonomy through human-AI collaboration’. We thank Mukul Sahu for his comments on FEM modeling.

REFERENCES

- [1] S. Kim, C. Laschi, and B. Trimmer, “Soft robotics: a bioinspired evolution in robotics,” *Trends in Biotechnology*, vol. 31, 2013.
- [2] D. Rus and M. Tolley, “Design, fabrication and control of soft robots,” *Nature*, vol. 521, 2015.
- [3] H. Yin, A. Varava, and D. Kragic, “Modeling, learning, perception, and control methods for deformable object manipulation,” *Science Robotics*, vol. 6, 2021.
- [4] N. El-Atab, R. Mishra, F. Al-modaf, L. Joharji, A. Alsharif, H. Alamoudi, M. Diaz, N. Qaiser, and M. Mustafa, “Soft actuators for soft robotic applications: A review,” *Advanced Intelligent Systems*, vol. 2, 2020.
- [5] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 5026–5033.
- [6] B. Amos, I. D. J. Rodriguez, J. Sacks, B. Boots, and J. Z. Kolter, “Differentiable MPC for end-to-end planning and control,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS)*, 2018, p. 8299–8310.
- [7] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018, pp. 1861–1870.
- [8] E. Coevoet, A. Escande, and C. Duriez, “Soft robots locomotion and manipulation control using fem simulation and quadratic programming,” in *Proceedings of 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, 2019, pp. 739–745.
- [9] T. Karner and J. Gotlih, “Position control of the dielectric elastomer actuator based on fractional derivatives in modelling and control,” *Actuators*, vol. 10, 2021.
- [10] T. Lopez-Guevara, N. K. Taylor, M. U. Gutmann, S. Ramamoorthy, and K. Subr, “Adaptable pouring: Teaching robots not to spill using fast but approximate fluid simulation,” in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 77–86.
- [11] S. David Müzel, E. Bonhin, N. Guimarães, and E. Guidi, “Application of the finite element method in the analysis of composite materials: A review,” *Polymers*, vol. 12, 2020.
- [12] U. Gupta, L. Qin, Y. Wang, H. Godaba, and J. Zhu, “Soft robots based on dielectric elastomer actuators: a review,” *Smart Materials and Structures*, vol. 28, 2019.
- [13] F. Zhou, X. Yang, Y. Xiao, Z. Zhu, T. Li, and Z. Xu, “Electromechanical analysis and simplified modeling of dielectric elastomer multilayer bending actuator,” *AIP Advances*, vol. 10, 2020.
- [14] J. Shintake, H. Shea, and D. Floreano, “Biomimetic underwater robots based on dielectric elastomer actuators,” in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 4957–4962.
- [15] M. Duduta, F. Berlinger, R. Nagpal, D. R. Clarke, R. J. Wood, and F. Z. Temel, “Tunable multi-modal locomotion in soft dielectric elastomer robots,” *IEEE Robotics and Automation Letters*, vol. 5, 2020.
- [16] J. Cao, W. Liang, Q. Ren, U. Gupta, F. Chen, and J. Zhu, “Modelling and control of a novel soft crawling robot based on a dielectric elastomer actuator,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 4188–4193.
- [17] O. Araromi and S. Burgess, “A finite element approach for modelling multilayer unimorph dielectric elastomer actuators with inhomogeneous layer geometry,” *Smart Materials and Structures*, vol. 21, 2012.
- [18] X. Zhao and Z. Suo, “Method to analyze programmable deformation of dielectric elastomer layers,” *Applied Physics Letters*, vol. 93, 2008.
- [19] P. Jamdagni and Y.-B. Jia, “Robotic cutting of solids based on fracture mechanics and fem,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 8252–8257.
- [20] E. Heiden, M. Macklin, Y. S. Narang, D. Fox, A. Garg, and F. Ramos, “DiSECT: A Differentiable Simulation Engine for Autonomous Robotic Cutting,” in *Proceedings of Robotics: Science and Systems*, 2021.
- [21] 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)*, 2020, pp. 417–423.
- [22] T. G. Thuruthel, E. Falotico, F. Renda, and C. Laschi, “Model-based reinforcement learning for closed-loop dynamic control of soft robotic manipulators,” *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 124–134, 2019.
- [23] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, 2019.
- [24] R. E. Pelrine, R. D. Kornbluh, and J. P. Joseph, “Electrostriction of polymer dielectrics with compliant electrodes as a means of actuation,” *Sensors and Actuators A: Physical*, vol. 64, 1998.
- [25] M. Smith, *ABAQUS/Standard User’s Manual, Version 6.9*. United States: Dassault Systèmes Simulia Corp, 2009.
- [26] H. Zhang, M. Dai, and Z. Zhang, “Application of viscoelasticity to nonlinear analyses of circular and spherical dielectric elastomers,” *AIP Advances*, vol. 9, no. 4, p. 045010, 2019.
- [27] V. Piefort, “Finite element modeling of piezoelectric structures,” 2000.
- [28] N. Wang, C. Chaoyu, H. Guo, B. Chen, and X. Zhang, “Advances in dielectric elastomer actuation technology,” *Science China Technological Sciences*, vol. 61, 2017.
- [29] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in PyTorch,” 2017.