Data–Driven Disturbance Observers for Estimating External Forces on Soft Robots

Cosimo Della Santina, Ryan Landon Truby, and Daniela Rus

***Abstract*—Unlike traditional robots, soft robots can intrinsically interact with their environment in a continuous, robust, and safe manner. These abilities - and the new opportunities they open - motivate the development of algorithms that provide reliable information on the nature of environmental interactions and, thereby, enable soft robots to reason on and properly react to external contact events. However, directly extracting such information with integrated sensors remains an arduous task that is further complicated by also needing to sense the soft robot's configuration. As an alternative to direct sensing, this paper addresses the challenge of estimating contact forces directly from the robot's posture. We propose a new technique that merges a nominal disturbance observer, a model-based component, with corrections learned from data. The result is an algorithm that is accurate yet sample efficient, and one that can reliably estimate external contact events with the environment. We prove the convergence of our proposed method analytically, and we demonstrate its performance with simulations and physical experiments.**

***Index Terms*—Modeling, control, and learning for soft robots, contact modeling, model learning for control.**

I. Introduction

SOFT-bodied robots can adaptively and safely interact with their environment due to their continuously deformable, compliant nature, opening up new opportunities and applications across the field of robotics [1]. However, it is soft robots' highly deformable nature that renders their control so challenging. Several algorithms and strategies for closed-loop control have been investigated in recent years [2]–[4], with most requiring some means of perceiving an estimation of the soft robot's configuration and discerning interactions with the external environment. While advances in soft sensor technologies have granted soft robots proprioceptive and tactile sensing capabilities [5], a complete and simultaneous sensing of body posture and contact events with the environment remains a grand challenge for the field.

Manuscript received February 24, 2020; accepted June 29, 2020. Date of publication July 21, 2020; date of current version July 28, 2020. This letter was recommended for publication by Associate Editor D. Shin and Editor C. Laschi upon evaluation of the reviewers' comments. This work was supported by the NSF EFRI Program under Grant 1830901. The work of Ryan Landon Truby was supported by the Schmidt Science Fellows program, in partnership with the Rhodes Trust. *(Cosimo Della Santina and Ryan Landon Truby contributed equally to this work.) (Corresponding author: Cosimo Della Santina.)*

The authors are with the MIT Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: cosimodellasantina@gmail.com; rltruby@mit.edu; rus@csail.mit.edu).

This letter has supplementary downloadable material available at https://ieeexplore.ieee.org, provided by the authors. The material contains a video overview of all contributions.

Digital Object Identifier 10.1109/LRA.2020.3010738

Rather than rely solely on body-integrated soft sensors, one alternative avenue for soft robotic perception of environmental interactions is to use algorithms that extract information from online readings of a partial set of readily accessible feedback motifs (e.g., motion capture data or actuation inputs), coupled with *a priori* knowledge of system behavior. Several works in the literature have taken this approach. For example, [6] shows that complete contact characterization in a soft fingertip can be extracted from a single force-torque sensor. In [7], a Kalman filter is used to integrate accelerations and angular velocities to postures, and in [8], the whole shape of a soft robot is regressed from the location of a finite number of markers attached to its body. Such algorithms are worth investigating, even for future embodiments of soft sensorized robots, as added algorithmic perception capabilities would only beneficially augment soft robot performance.

In this paper, we focus on the challenge of characterizing contact events on a soft robot when only knowledge of its posture and actuation are available. This problem has been considered in simulation for continuum robots with high Young's modulus in [9]. In [10], the sole contact location of a single pneumatic bending actuator is estimated by inverting the kinematic model of the system. More recently, in [11], a long short-term memory (LSTM) network is used to learn the relationship connecting pressure inputs and soft strain sensor readings to the forces applied by the tip of a soft finger. While these works present exciting results, they have practical limitations. Namely, they deal with simple systems or include very limited experimental validation.

To overcome these limitations, we present an algorithm combining model-based techniques with machine learning. Central to our approach is a nonlinear disturbance observer, a model-based component designed to achieve complete and exact knowledge of external forces under Piecewise Constant Curvature (PCC) assumptions [4] (see [12], [13] for examples of disturbance observers for rigid bodied robots). However, in the general case, the real system will present some discrepancies with respect to (w.r.t.) this idealized view. We capture this behavior through a stochastic model, and we use machine learning to learn the model from the data via Gaussian Process Regressors (GPRs) [14]. This data-driven component is integrated into our algorithm. Convergence is proven analytically in the uncertain and stochastic case, as well as experimentally on a soft robot arm moving in 3D space via 12 discrete fluidic actuation networks (i.e., 4 in each segment) [15]. This work contributes:

1) A new class of disturbance observers able to estimate contact forces in soft robots modelled with PCC,
2) An algorithm that integrates models of uncertainty learned from data, using the proposed model-based disturbance observers,

3) Proofs of convergence in nominal and uncertain conditions, and
4) quantitative simulations and qualitative experiments.

It is worth underscoring that in this paper we are interested in capturing low frequency contact information, and we do not claim that our method is appropriate for instantaneous reaction to contact events [13]. Still, important possible applications enabled by our method include kinesthetic teaching [16], mapping environment geometry through touch [17], and environmental stiffness sensing [18].

II. Background: Dynamic 3D PCC Model

A serial PCC robot is a mechanical system composed of a sequence of continuously deformable segments. Their curvature is constant in space (CC) and variable in time, and they are connected so that the resulting curve is everywhere differentiable (see an example in Fig. 2(a)). PCC models have proven to be a good finite-dimensional description of soft robots under a variety of conditions and for applications ranging from design to control [4], [19].

To describe the robot's configuration, we introduce n reference frames, $\{S_1\},\dots,\{S_n\}$, attached at the ends of each segment, plus one fixed base frame, $\{S_0\}$. We consider four arcs for the $i-$th segment, as depicted in Fig. 2(b). The first arc has one end connected to $[d,0,0]^{\mathrm{T}}$ expressed in $\{S_{i-1}\}$, and the other to $[d,0,0]^{\mathrm{T}}$ expressed in $\{S_i\}$. $d\in\mathbb{R}^+$ is the radius of the cross-section of the segment, assumed here without loss of generality (w.l.o.g.) to be circular and constant along the robot. Similarly, the other three arcs are connected to $[-d,0,0]^{\mathrm{T}}$, $[0,d,0]^{\mathrm{T}}$, and $[0,-d,0]^{\mathrm{T}}$, respectively. We call $L_{i,1},L_{i,2},L_{i,3},L_{i,4}\in\mathbb{R}$ the lengths of these arcs. These quantities can be linearly combined in the following manner: $\Delta_{\mathrm{x},i}=\frac{L_{2,i}-L_{1,i}}{2}$, $\Delta_{\mathrm{y},i}=\frac{L_{4,i}-L_{3,i}}{2}$. Defined this way, $\Delta_{\mathrm{x},i}$ and $\Delta_{\mathrm{y},i}$ provide a complete parametrization of the configuration manifold of the $i-$th CC segment. The formal proof of this statement is provided in [20]. We call T^i_{i-1} the homogeneous transformations mapping $\{S_i\}$ into $\{S_{i-1}\}$, i.e. connecting the configuration $\Delta_{\mathrm{x},i},\Delta_{\mathrm{y},i}$ to the posture of $\{S_i\}$, expressed in coordinates $\{S_{i-1}\}$. The matrix representation of this SE(3) element is

$$T^i_{i-1}(\Delta_{\mathrm{x},i},\Delta_{\mathrm{y},i})=\begin{bmatrix}R^i_{i-1}(\Delta_{\mathrm{x},i},\Delta_{\mathrm{y},i}) & t^i_{i-1}(\Delta_{\mathrm{x},i},\Delta_{\mathrm{y},i})\\ [0\quad 0\quad 0] & 1\end{bmatrix},$$

with

$$R^i_{i-1}=\begin{bmatrix}1+\frac{\Delta^2_{\mathrm{x},i}}{\Delta^2}\mathrm{C}_i(\Delta_i) & \frac{\Delta_{\mathrm{x},i}\Delta_{\mathrm{y},i}}{\Delta^2}\mathrm{C}_i(\Delta_i) & -\frac{\Delta_{\mathrm{x},i}}{\Delta}\mathrm{S}_i(\Delta_i)\\ \frac{\Delta_{\mathrm{x},i}\Delta_{\mathrm{y},i}}{\Delta^2}\mathrm{C}_i(\Delta_i) & 1+\frac{\Delta^2_{\mathrm{y},i}}{\Delta^2}\mathrm{C}_i(\Delta_i) & -\frac{\Delta_{\mathrm{y},i}}{\Delta}\mathrm{S}_i(\Delta_i)\\ \frac{\Delta_{\mathrm{x},i}}{\Delta}\mathrm{S}_i(\Delta_i) & \frac{\Delta_{\mathrm{y},i}}{\Delta}\mathrm{S}_i(\Delta_i) & 1+\mathrm{C}_i(\Delta_i)\end{bmatrix},$$

$$t^i_{i-1}=\frac{d_iL_{0,i}}{\Delta^2}\begin{bmatrix}\mathrm{C}_i(\Delta_i)\Delta_{\mathrm{x},i} & \mathrm{C}_i(\Delta_i)\Delta_{\mathrm{y},i} & \mathrm{S}_i(\Delta_i)\Delta_i\end{bmatrix}^{\mathrm{T}},$$

where $\Delta_i=\sqrt{\Delta^2_{\mathrm{x},i}+\Delta^2_{\mathrm{y},i}}$, $\mathrm{S}_i(\Delta_i)=\sin(\frac{\Delta}{d_i})$, $\mathrm{C}_i(\Delta_i)=(\cos(\frac{\Delta}{d_i})-1)$. With the kinematics of the robot fully specified, the dynamics can be derived, for example, by Lagrangian formulation or by matching the dynamics of the PCC robot with the one of a constrained rigid robot [4]. The resulting general form of the dynamics is

$$B(q)\ddot{q}+C(q,\dot{q})\dot{q}+Kq+D\dot{q}+G(q)=A(q)u+\tau_{\mathrm{ext}},\quad(1)$$

where $q\in\mathbb{R}^{2n}$ is the configuration vector containing $\Delta_{\mathrm{x},i}$ and $\Delta_{\mathrm{y},i}$ for all segments, and $\dot{q},\ddot{q}\in\mathbb{R}^{2n}$ are its time derivatives. $B(q)\in\mathbb{R}^{2n\times 2n}$ is the inertia matrix, $C(q,\dot{q})\dot{q}\in\mathbb{R}^{2n}$ contains the Coriolis and centrifugal forces. $K\in\mathbb{R}^{2n}$ is the stiffness, and $D\in\mathbb{R}^{2n}$ is the damping. $G(q)\in\mathbb{R}^{2n}$ is the gravitational force acting on the robot. $A(q)\in\mathbb{R}^{2n\times m}$ maps the input $u\in\mathbb{R}^m$ in wrenches producing independent accelerations $\ddot{q}$. Finally, $\tau_{\mathrm{ext}}\in\mathbb{R}^{2n}$ is an external torque produced by some interaction with an unstructured environment.

III. Disturbance Observers for Piecewise Constant Curvature Robots

The class of algorithms we present, which we call *nominal disturbance observers of PCC robots*, can accurately estimate external torques τ_{ext} from the knowledge of the robot's configuration q and actuation u. Apart from the obvious differences in terms of modeling the robot, the challenge of estimating external torques in this manner has some specificity w.r.t. the rigid bodied case [12], [13], even when considering a PCC approximation. First, we want to be independent here from the knowledge of $\dot{q}$, since it is a quantity that, while being well defined, is very hard to measure in soft robots. We also do not want to assume the knowledge of B and C, which can be unreliable for coarse space discretizations. To this end, we assume small accelerations (i.e. $B(q)\ddot{q}\simeq 0$) and that damping forces are dominant w.r.t. Coriolis and centrifugal ones (i.e. $||D\dot{q}||>>||C(q,\dot{q})\dot{q}||$). Under these assumptions, (1) can be approximated with the following set of n first order nonlinear differential equations

$$D\dot{q}=-Kq-G(q)+A(q)u+\tau_{\mathrm{ext}},\quad(2)$$

with elements defined as in (1). It is very important to underline that this approximation is instrumental to deriving the algorithms. However, convergence results will be provided for the full system (1).

A. Algorithm

The estimation problem is resolved by integrating the following set of ordinary differential equations, which describe the dynamics of a class of nonlinear disturbance observers for PCC robots:

$$\begin{aligned}D\dot{\hat{q}}&=-Kq-G(q)+A(q)u+\left(1+\sum_{x=\eta,u,q}\frac{\nabla_x\gamma}{\gamma^2}\dot{x}\right)\hat{\tau}_{\mathrm{ext}},\\ \dot{\eta}&=f(\eta,q,u),\\ \hat{\tau}_{\mathrm{ext}}&=\gamma(\eta,u,q)D\left(q-\hat{q}\right),\end{aligned}\quad(3)$$

where the role of the correction $\sum_{x=\eta,u,q}\frac{\nabla_x\gamma}{\gamma^2}\dot{x}$ will become clear later. $\eta\in\mathbb{R}^k$ is an internal state with dynamics described by the smooth function $f:\mathbb{R}^k\times\mathbb{R}^m\times\mathbb{R}^{2n}\to\mathbb{R}^k$. $\gamma:\mathbb{R}^k\times\mathbb{R}^m\times\mathbb{R}^{2n}\to\mathbb{R}$ is a generic (possibly non-smooth) function. $\hat{q}$ is an internal representation of the configuration q. All the other terms are defined as in (1). Note that q here is an input to the algorithm, not a state variable.

Our algorithm is depicted as a block scheme in Fig. 3, with the structure of (1) and (3) highlighted in green and blue, respectively. The first set of differential equations in (3) is the core of the proposed disturbance observer. This is an internal replica of the robot's dynamics (2). This part of the observer is largely specified by the knowledge of the model, and it is what performs the most important computations for assuring the algorithm functions correctly. The second set of differential equations in (3) is used to estimate variables of interest in the function γ. Any choice of f defining a proper set of ordinary differential equations is admitted. Finally, the third set of equations in (3) is

algebraic and serves as an output function, where γ should be regarded as the instantaneous convergence rate of the observer. Therefore, f and γ can be chosen so to implement heuristics speeding up or slowing down the algorithm's convergence under specific conditions.

The following theorems provide theoretical support to these considerations.

Theorem 1: Consider the system (3) fed with q measured from (2). If $\dot{\tau}_{\text{ext}} = 0$ and $\int_0^t \gamma \to \infty$ for $t \to \infty$, then $\hat{\tau}_{\text{ext}} \to \tau_{\text{ext}}$.

Proof: First, we evaluate the time derivative of the estimation:

$$\begin{aligned}\dot{\hat{\tau}}_{\text{ext}} &= \gamma(D\dot{q} - D\dot{\hat{q}}) + \dot{\gamma}D(q - \hat{q}) \\ &= \gamma\left(\tau_{\text{ext}} - \left(1 + \sum_{x=\eta,u,q} \frac{\nabla_x \gamma}{\gamma^2}\dot{x}\right)\hat{\tau}_{\text{ext}}\right) \\ &\quad + \dot{\gamma}D(q - \hat{q}). \end{aligned} \tag{4}$$

Using the chain rule, the time derivative of γ can be expressed as $\dot{\gamma} = \nabla_\eta \gamma\, \dot{\eta} + \nabla_u \gamma\, \dot{u} + \nabla_q \gamma\, \dot{q}$, and (4) becomes

$$\begin{aligned}\dot{\hat{\tau}}_{\text{ext}} &= \gamma\left(\tau_{\text{ext}} - \left(1 + \frac{\dot{\gamma}}{\gamma^2}\right)\hat{\tau}_{\text{ext}}\right) + \dot{\gamma}D(q - \hat{q}) \\ &= \gamma(\eta, u, q)(\tau_{\text{ext}} - \hat{\tau}_{\text{ext}}). \end{aligned} \tag{5}$$

The last step is true since $\hat{\tau}_{\text{ext}} = \gamma(\eta, u, q)D(q - \hat{q})$. This is a set of decoupled first order linear differential equations, with time variant gains γ. The evolution is [21]

$$\begin{aligned}\hat{\tau}_{\text{ext}}(t) &= e^{-\int_0^t \gamma}\left(\hat{\tau}_{\text{ext}}(0) + \int_0^t \gamma \tau_{\text{ext}} e^{\int_0^{\bar{t}} \gamma} \mathrm{d}\bar{t}\right) \\ &= e^{-\int_0^t \gamma}\left(\hat{\tau}_{\text{ext}}(0) + \tau_{\text{ext}} \int_0^t \gamma e^{\int_0^{\bar{t}} \gamma} \mathrm{d}\bar{t}\right). \end{aligned} \tag{6}$$

Since $\frac{\mathrm{d}e^{\int_0^t \gamma}}{\mathrm{d}t} = \gamma e^{\int_0^t \gamma}$, we have the closed form solution

$$\begin{aligned}\hat{\tau}_{\text{ext}}(t) &= e^{-\int_0^t \gamma}\left(\hat{\tau}_{\text{ext}}(0) + \tau_{\text{ext}}\left(e^{\int_0^t \gamma} - \mathbf{1}\right)\right) \\ &= e^{-\int_0^t \gamma}\hat{\tau}_{\text{ext}}(0) + (\mathbf{1} - e^{-\int_0^t \gamma})\tau_{\text{ext}} \to \tau_{\text{ext}}, \end{aligned} \tag{7}$$

where $\mathbf{1}$ is a vector of the opportune length with all elements equal to 1. The last step holds under the hypothesis $\int_0^t \gamma \to \infty$ for $t \to \infty$. ■

Theorem 1 can be further extended to assure the convergence of the algorithm under fully dynamic conditions. Note that no model of the dynamic forces is added to the observer itself.

Theorem 2: Consider the system (3) fed with data extracted from (1). If the same hypotheses of Theorem 1 are fulfilled and the rate of exponential convergence α of $||B(q)\ddot{q} + C(q, \dot{q})\dot{q}||$ to 0, is such that $\alpha > \max_t \gamma - \lim_{t\to\infty} \frac{\int_0^t \gamma}{t}$, then $\hat{\tau}_{\text{ext}} \to \tau_{\text{ext}}$.

Proof: The same steps of the proof of Theorem 1 can be repeated until (7), which becomes

$$\begin{aligned}\hat{\tau}_{\text{ext}}(t) &= e^{-\int_0^t \gamma}\hat{\tau}_{\text{ext}}(0) + (1 - e^{-\int_0^t \gamma})\tau_{\text{ext}} \\ &\quad - e^{-\int_0^t \gamma}\int_0^t \gamma(s^+(\bar{t}) + s^-(\bar{t}))e^{\int_0^{\bar{t}} \gamma}\mathrm{d}\bar{t}, \end{aligned} \tag{8}$$

where

$$\begin{aligned} s^+ &= \begin{cases} B(q)\ddot{q} + C(q, \dot{q})\dot{q} & \text{if } B(q)\ddot{q} + C(q, \dot{q})\dot{q} > 0 \\ 0 & \text{otherwise,} \end{cases} \\ s^- &= \begin{cases} B(q)\ddot{q} + C(q, \dot{q})\dot{q} & \text{if } B(q)\ddot{q} + C(q, \dot{q})\dot{q} < 0 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{9}$$

The exponential convergence of $||B(q)\ddot{q} + C(q, \dot{q})\dot{q}||$ implies that $S, \alpha > 0$ exist such that $0 < s^+(t) < Se^{-\alpha t}$, $-Se^{-\alpha t} < s^-(t) < 0$, $\forall t$. We can bound the third term of (8) as follows:

$$\begin{aligned} &||e^{-\int_0^t \gamma}\int_0^t \gamma(s^+ + s^-)e^{\int_0^{\bar{t}} \gamma}\mathrm{d}\bar{t}|| \\ &\leq 2S||e^{-\int_0^t \gamma}\int_0^t \gamma e^{-\alpha\bar{t} + \int_0^{\bar{t}} \gamma}\mathrm{d}\bar{t}|| \\ &= 2S||e^{-\int_0^t \gamma}\left(e^{-\alpha t + \int_0^t \gamma} - \mathbf{1} + \alpha\int_0^t e^{-\alpha\bar{t} + \int_0^{\bar{t}} \gamma}\mathrm{d}\bar{t}\right)|| \\ &\leq 2S||e^{-\alpha t} - e^{-\int_0^t \gamma} + e^{-\int_0^t \gamma}\alpha\int_0^t e^{-(\alpha - \max\gamma)\bar{t}}\mathrm{d}\bar{t}|| \\ &= 2S||e^{-\alpha t} - e^{-\int_0^t \gamma} - \frac{\alpha}{\alpha - \gamma^+}e^{-\alpha t + ((\max\gamma)t - \int_0^t \gamma)}|| \to 0, \end{aligned} \tag{10}$$

where we exploit the upper bound on α assumed by hypothesis, to say that $e^{-\alpha t + ((\max\gamma)t - \int_0^t \gamma)} \to 0$.

This implies that the argument goes to zero. Combining (7), (8), and (10) concludes the proof. ■

Note that it is typically the case that $\tau_{\text{ext}}, u, q(0)$, and $\dot{q}(0)$ are such that $\ddot{q}, \dot{q}$ converge exponentially to 0 due to the self-stabilizing properties of soft robots [22]. Under such conditions, it is always true that $||B(q)\ddot{q} + C(q, \dot{q})\dot{q}||$ converges exponentially to 0. Indeed, $||B(q)||$ can always be limited with an affine function of $||q||$, and $||C(q, \dot{q})\dot{q}||$ with a quadratic function of $||q||$ and $||\dot{q}||$.

B. Two Examples of Observers

Here, we consider two implementations of the general class of observes introduced in the previous section. First, we take a constant γ, obtaining

$$\begin{aligned} D\dot{\hat{q}} &= -Kq - G(q) + A(q)u + \hat{\tau}_{\text{ext}}, \\ \hat{\tau}_{\text{ext}} &= \gamma_0 D(q - \hat{q}). \end{aligned} \tag{11}$$

Note that for this observer, having the soft robot reach a steady state is sufficient for also achieving convergence in the observer. Indeed, $\max_t \gamma_0 = \gamma_0$ and $\lim_{t\to\infty} \frac{\int_0^t \gamma_0}{t} = \gamma_0$, leading to the simplified condition $\alpha > 0$ in Theorem 2.

As an example of a non-constant γ, we desensitize the algorithm to dynamic forces detected during transients. This can be done using the following observer:

$$\begin{aligned} D\dot{\hat{q}} &= -Kq - G(q) + A(q)u \\ &\quad + \left(1 + \frac{10^2\gamma_0\eta_3^{\mathrm{T}}(p^3(q - \eta_1) - p^2\eta_2 - p\eta_3)}{(1 + ||10\eta_3||^2)^2}\right)\hat{\tau}_{\text{ext}}, \\ \begin{bmatrix}\dot{\eta}_1 \\ \dot{\eta}_2 \\ \dot{\eta}_3\end{bmatrix} &= \begin{bmatrix}-pI & 0 & 0 \\ -p^2 I & -pI & 0 \\ -p^3 I & -p^2 I & -pI\end{bmatrix}\begin{bmatrix}\eta_1 \\ \eta_2 \\ \eta_3\end{bmatrix} + \begin{bmatrix}pI \\ p^2 I \\ p^3 I\end{bmatrix} q, \\ \hat{\tau}_{\text{ext}} &= \frac{1}{2}\frac{\gamma_0}{1 + ||10\eta_3||^2}D(q - \hat{q}). \end{aligned} \tag{12}$$

Note that η_3 is a low pass approximation of $\ddot{q}$. p defines the speed of convergence of this estimation and is taken here to be equal to 10. I is the identity matrix. The reasoning behind this choice of f and γ is that when $\ddot{q} \simeq 0$, the effect of $B(q)\ddot{q}$ will be negligible, even if the robot's inertia is large. γ_0 defines the

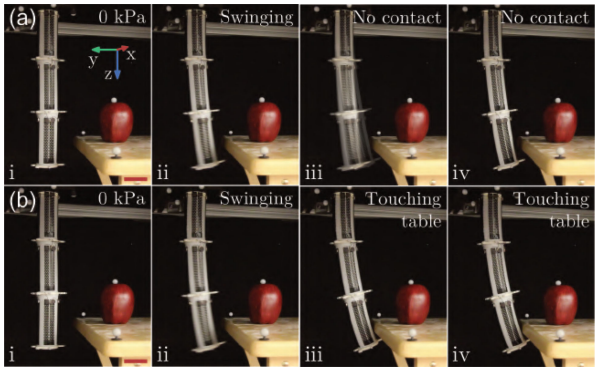

Fig. 1. Contact Perception in Soft Robots. (a, b) Images from Experiment 1 (discussed herein) show our soft robot arm actuating such that it fails (a) and succeeds (b) in contacting a table edge. In these tests, the soft robot is actuated from rest (i) towards the table (ii) and held at constant actuation pressure (iii) until a steady-state configuration is reached (iv). Our control system estimates these contact events using only posture information and actuation pressures. (Reference frame provided in (a, i); scale bars = 5 cm).

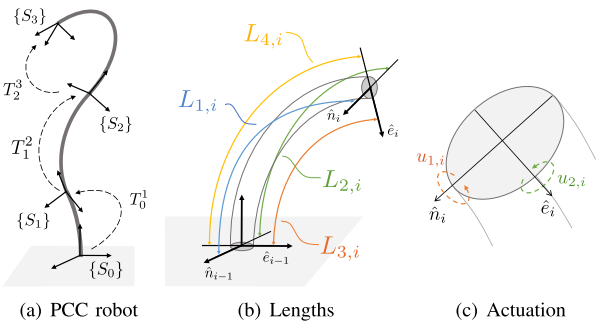

(a) PCC robot (b) Lengths (c) Actuation

Fig. 2. PCC Soft Robots. PCC models are an effective way of approximating soft robots by using a reduced set of state variables. Panel (a) shows an example of a PCC robot with three segments. Reference systems, $\{S_i\}$, and transformations between them, T_{i-1}^i, are also shown. Panel (b) shows a single segment with the four arcs used to define the robot's configuration space. Panel (c) shows the actuation that we consider in the validation. Note, however, that the proposed method is not limited to this type of actuation.

rate of convergence when $\ddot{q}$ is small. There are other possible cases of f and γ that we do not investigate here for the sake of space: decreasing γ when an estimation of $||B(q)\ddot{q} + C(q,\dot{q})\dot{q}||$ is large, and increasing γ when $\dot{u} \simeq 0$ for a sufficiently long period of time. These values can either be estimated using dirty derivatives as employed in (12), though direct measurements, or using model based observers.

C. Simulations

We simulate a PCC robot with three segments serially connected. Each segment is 0.1m long. The robot's weight (0.1 kg/segment) is homogeneously distributed along its length. We consider it actuated through a pair of internal torques acting around $\hat{n}_i$ and $\hat{e}_i$ (see Fig. 2(c)). The resulting input matrix $A(q)$ is block diagonal, with diagonal blocks $A_i(\Delta_{\mathrm{x},i}, \Delta_{\mathrm{y},i}) = J_i^{\mathrm{T}} R_0^i - J_{i-1}^{\mathrm{T}} R_0^{i-1}$, where J_i is the Jacobian matrix mapping $\dot{\Delta}_{\mathrm{x},i}, \dot{\Delta}_{\mathrm{y},i}$ into the angular velocity expressed around $\hat{n}_i$ and $\hat{e}_i$. R_0^i is the rotation matrix mapping $\{S_i\}$ into $\{S_0\}$. The impedance is $K = 3I$ N, and the damping is $D = 3I$ Ns. The gravity field is directed downward, perpendicular to the robot's base. The robot is simulated using the full model (1). The initial configuration, $q(0) = 0$, is straight, as shown in Fig. 1 b. The state of the observer is initialized in $\hat{q} = 0$, $\eta = 0$. We evaluate both the (11) and (12) cases, with $\gamma_0 = 10$. An input u is generated at $t = 0$. It lasts until the end of the simulation. An external force τ_{ext} is applied at $t = 15$s. We present two simulations that differ in the choices of u and τ_{ext}.

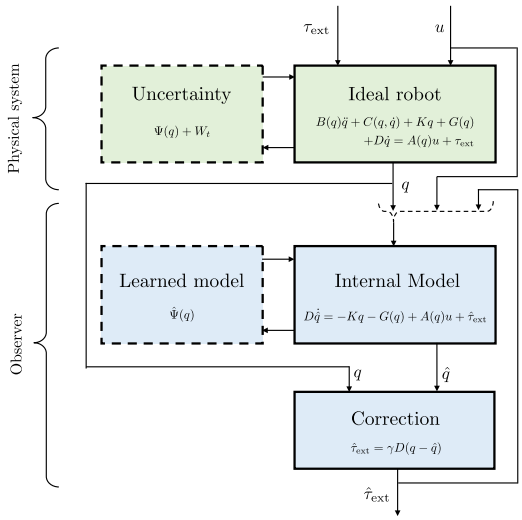

Fig. 3. General Scheme of the Algorithm, for $\dot{\gamma} = 0$. Green and blue blocks represent the real system and proposed observer, respectively. The right column of blocks represent those pertaining to the nominal case, while the left column (also shown with dashed edges) corresponds to those generalizing to the generic soft robotic case.

In the first simulation, u and τ_{ext} are step functions with steady state $[0.05, -0.05]$Nm and $[-0.25, 0.25]$N, respectively, for each segment. Evolution of the robot's configuration is shown in Fig. 4(a). The behaviors of the two observers (11) and (12) are reported in Fig. 4(b) and (c), respectively. Both observers reach the correct steady state, but (12) is more effective in rejecting the cross-talk from dynamic forces. The simulation was repeated 100 times with τ_{ext} extracted from a uniform distribution with an upper bound of 0.3Nm and lower bound of -0.3Nm. The steady state error was always zero.

In the second simulation, a step force is produced in Cartesian coordinates. The application point is the middle of Segment 2. Its value in global coordinates is $[1\ 0\ 0]^{\mathrm{T}}$N. The input u is the same as Simulation 1. Fig. 4(d) shows the evolution of the configuration in time. Fig. 4(e) and (f) depict the behavior of the two observers. Again, (12) with variable γ rejects oscillations induced by dynamic forces more effectively. This simulation was repeated 100 times. The amplitude of the external force was extracted from a uniform distribution with upper bound 1N and lower bound -1N. Again, the steady state error was always zero.

IV. Bridging the Reality Gap Through Machine Learning

Results presented so far hold under the assumption that the soft robot can be exactly modeled as a PCC system. However, this assumption is not always accurate in practice, especially if the discretization is coarse. Note that a finer discretization implies a higher number of sensor readings. We consider here

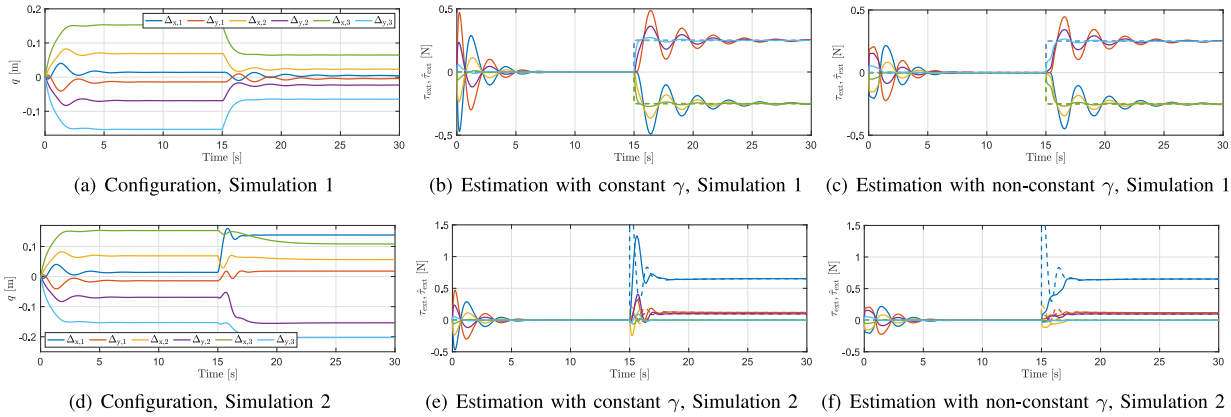

(a) Configuration, Simulation 1 (b) Estimation with constant γ, Simulation 1 (c) Estimation with non-constant γ, Simulation 1

(d) Configuration, Simulation 2 (e) Estimation with constant γ, Simulation 2 (f) Estimation with non-constant γ, Simulation 2

Fig. 4. Simulation Results. u and τ_{ext} are step functions, applied starting at $t = 0$s and $t = 15$s, respectively. In Simulation 1 (a-c), τ_{ext} is constant in configuration space; in Simulation 2 (d-f), it is constant in Cartesian coordinates. (a),(d) show the evolution of robot's configuration. (b),(e) report the estimated $\hat{\tau}_{\text{ext}}$ (solid lines) compared to the force to be estimated τ_{ext} (dashed lines) when observer (11) is used. (c),(f) correspond to when (12) is used.

the use of machine learning to compensate for model mismatches. We start by representing un-modeled effects through an unknown function of the robot's configuration $\Psi : \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$, plus a stochastic uncertainty in the form of a continuous time white noise W_t, with null mean and co-variance function $\mathrm{E}[W_{t_1} W_{t_2}] = \sigma^2 \delta(t_1 - t_2)$, where δ is the Kronecker delta function. Their $i-$th components are Ψ_i and $W_{t,i}$. Thus, (2) is modified into the stochastic differential equation [23]

$$D\dot{q} = -Kq - G(q) + A(q)u + \tau_{\text{ext}} + \Psi(q) + W_t, \quad (13)$$

with main quantities defined as in (1).

A. Algorithm

The goal of the new algorithm is to obtain an estimation of the disturbance τ_{ext} under the uncertain conditions described by (13). To do this, we use machine learning to regress a suitable approximation of $\Psi(q)$ from data. We then augment the model-based observers proposed in the previous section to include the learned model (see Fig. 3). For the sake of space, we focus here on the constant gain case, i.e. $\dot{\gamma} = 0$.

At any time, the algorithm may or may not be in the training phase. The only supervision that the algorithm needs is to know in which phase it is. In both phases, the following computations are continuously performed

$$\begin{aligned} D\dot{\hat{q}} &= -Kq - G(q) + A(q)u + \hat{\Psi}(q) + \hat{\tau}_{\text{ext}}, \\ \hat{\tau}_{\text{ext}} &= \gamma D(\hat{q} - q), \end{aligned} \quad (14)$$

where $\hat{\Psi}$ is a suitable estimation of Ψ, with $i-$th component called $\hat{\Psi}_i$. The remaining elements are as in (3). If the algorithm is in the training phase, it will assume that no contacts with the environment are happening. So, the following can be assumed anytime the observer (14) reaches a steady state: $\hat{\tau}_{\text{ext}} \simeq Kq + G(q) - A(q)u - \hat{\Psi}(q) = \Psi(q) - \hat{\Psi}(q) + W_t$, where we substituted (13) evaluated in $\tau_{\text{ext}} = 0$ and $\dot{q} = 0$. Using these equations, examples $\{q, \Psi(q)\}$, which are not available since Ψ is unknown, can be generated as $\{q, \hat{\tau}_{\text{ext}} + \hat{\Psi}(q)\}$. The dataset composed by these pairs is then used to update the estimation $\hat{\Psi}$. While several machine learning techniques could be applied here, we use Gaussian process regressors (GPRs, see [14] for more details on this technique). GPRs inherently describe stochastic processes with Gaussian distribution. They are also proven to be the optimal interpolator under the considered hypotheses. Lastly, GPRs are very sample efficient, which is one of the main goals of our architecture. The result is Algorithm 1.

Algorithm 1: High Level Execution Flow.

```
while True do
    [q, u] ← ReadSensors()
    Ψ̂ ← EvaluateGPR(q, trainSet)
    [τ̂_ext, q̂̇] ← EvaluateObserver(q, u, Ψ̂)
    if ||q̂̇|| < ε && TrainingPhase() == True then
        trainSet ← Append(trainSet, {q, τ̂_ext + Ψ̂})
    end if
end while
```

The final theorem assures that the proposed observer converges in the neighborhood of the exact value, with a steady-state error given by the accuracy of $\hat{\Psi}$.

Theorem 3: Consider the system (14) fed with data extracted from (13). Let $||\Psi(q) - \hat{\Psi}(q)||$ be lower than Δ, $\forall q$.

If $\dot{\tau}_{\text{ext}} = 0$ and $\gamma > 0$, then a $\mu \in \mathbb{R}^+$ exists such that

$$||\mathrm{E}[\hat{\tau}_{\text{ext}}] - \tau_{\text{ext}}|| \rightarrow \mu < \Delta, \quad \mathrm{Var}[\hat{\tau}_{\text{ext},i}] \rightarrow \frac{\gamma}{2}\sigma^2, \quad (15)$$

with rate of exponential convergence γ.

Proof: Following the same steps of the proof of Theorem 1, we get $\dot{\hat{\tau}}_{\text{ext}} = -\gamma\hat{\tau}_{\text{ext}} + \gamma\tau_{\text{ext}} + s(t) + \gamma W_t$, where $s(t) = \gamma(\Psi(q(t)) - \hat{\Psi}(q(t)))$. Since this is a linear stochastic differential equation, its average evolution $\mathrm{E}[\hat{\tau}_{\text{ext}}]$ can be integrated in closed form [23] as

$$\begin{aligned} &\mathrm{E}\left[e^{-\gamma t}\hat{\tau}_{\text{ext}}(0) + (\mathbf{1} - e^{-\gamma t})\tau_{\text{ext}} + \int_0^t e^{-\gamma(t-\bar{t})}\left(\gamma W_{\bar{t}} + s(\bar{t})\right)\mathrm{d}\bar{t}\right] \\ &= e^{-\gamma t}\hat{\tau}_{\text{ext}}(0) + (\mathbf{1} - e^{-\gamma t})\tau_{\text{ext}} + \int_0^t e^{-\gamma(t-\bar{t})}(\gamma\mathrm{E}[W_{\bar{t}}] + s(\bar{t}))\mathrm{d}\bar{t} \\ &= e^{-\gamma t}\hat{\tau}_{\text{ext}}(0) + (\mathbf{1} - e^{-\gamma t})\tau_{\text{ext}} + \int_0^t e^{-\gamma(t-\bar{t})}s(\bar{t})\mathrm{d}\bar{t}, \end{aligned} \quad (16)$$

where we use the linearity of the integral and the hypothesis that $\mathrm{E}[W_{\bar{t}}] = 0$. Subtracting τ_{ext} from both sides of the equation and

extracting the norm yields

$$
\begin{aligned}
&||\mathrm{E}[\hat{\tau}_{\mathrm{ext}}](t) - \tau_{\mathrm{ext}}|| \\
&\leq ||e^{-\gamma t}(\hat{\tau}_{\mathrm{ext}}(0) + \tau_{\mathrm{ext}})|| + \int_0^t ||e^{-\gamma(t-\bar{t})} s(\bar{t})|| \mathrm{d}\bar{t} \\
&\leq ||e^{-\gamma t}(\hat{\tau}_{\mathrm{ext}}(0) + \tau_{\mathrm{ext}})|| + \int_0^t e^{-\gamma(t-\bar{t})} \max_{\bar{t}} ||s(\bar{t})|| \mathrm{d}\bar{t} \\
&\leq ||e^{-\gamma t}(\hat{\tau}_{\mathrm{ext}}(0) + \tau_{\mathrm{ext}})|| + (\mathbf{1} - e^{-\gamma t})\Delta \to \Delta \,,
\end{aligned}
$$

where we exploit the triangular inequality and the linearity of the integral. This concludes the first part of the proof.

We consider now the variance of the estimation

$$
\begin{aligned}
&\mathrm{Var}[\hat{\tau}_{\mathrm{ext},i}] = -\mathrm{E}[\hat{\tau}_{\mathrm{ext},i}]^2 + \mathrm{E}[\hat{\tau}^2_{\mathrm{ext},i}] \\
&= -\mathrm{E}[\hat{\tau}_{\mathrm{ext},i}]^2 + \mathrm{E}\left[\left(\int_0^t e^{-\gamma(t-\bar{t})} \gamma W_{\bar{t},i} \mathrm{d}\bar{t}\right)^2\right] \\
&\quad + 2\left(e^{-\gamma t}\hat{\tau}_{\mathrm{ext},i}(0) + (1 - e^{-\gamma t})\tau_{\mathrm{ext},i}\right. \\
&\quad \left. + \int_0^t e^{-\gamma(t-\bar{t})} s_i(\bar{t}) \mathrm{d}\bar{t}\right)\left(\int_0^t e^{-\gamma(t-\bar{t})} \gamma \mathrm{E}[W_{\bar{t},i}] \mathrm{d}\bar{t}\right) \\
&\quad + \left(e^{-\gamma t}\hat{\tau}_{\mathrm{ext},i}(0) + (1 - e^{-\gamma t})\tau_{\mathrm{ext},i} + \int_0^t e^{-\gamma(t-\bar{t})} s_i(\bar{t}) \mathrm{d}\bar{t}\right)^2 \\
&= \mathrm{E}\left[\left(\int_0^t e^{-\gamma(t-\bar{t})} \gamma W_{\bar{t},i} \mathrm{d}\bar{t}\right)^2\right],
\end{aligned}
$$

where we made $\hat{\tau}^2_{\mathrm{ext},i}$ explicit, exploited the linearity of the operators involved, and use (16). We now exploit the square of the integral as double integral

$$
\begin{aligned}
&\mathrm{E}\left[\int_0^t \int_0^t e^{-\gamma(t-\bar{t})} e^{-\gamma(t-\tilde{t})} \gamma^2 W_{\bar{t},i} W_{\tilde{t},i} \mathrm{d}\bar{t}\mathrm{d}\tilde{t}\right] \\
&= \int_0^t \int_0^t e^{-\gamma(t-\bar{t})} e^{-\gamma(t-\tilde{t})} \gamma^2 \mathrm{E}\left[W_{\bar{t},i} W_{\tilde{t},i}\right] \mathrm{d}\bar{t}\mathrm{d}\tilde{t} \\
&= \gamma^2\sigma^2 \int_0^t \int_0^t e^{-\gamma(2t-\bar{t}-\tilde{t})} \delta(\bar{t} - \tilde{t}) \mathrm{d}\bar{t}\mathrm{d}\tilde{t} \\
&= \gamma^2\sigma^2 \int_0^t e^{-\gamma(2t-2\bar{t})} \mathrm{d}\bar{t} = \frac{\gamma}{2}\sigma^2(1 - e^{-2\gamma t}) \to \frac{\gamma}{2}\sigma^2
\end{aligned}
$$

■

B. Experiments

To confirm our simulations work on real systems, we perform contact experiments on a soft robotic arm from [24] that is based on the designs introduced in [15]. This 12-actuator system is made of three silicone segments, and each can be independently inflated through four discrete fluidic chambers to produce an actuation pattern analogous to those considered in simulation. The weight of each segment is 135g, each segment's length is 109mm, and the radius of each segment cross-section is 20mm. The stiffness is assumed constant and estimated to be $4.241 \frac{\mathrm{N}}{\mathrm{mm}}$. The damping is estimated to be the matrix with diagonal elements $[.48,\ .65,\ .48,\ .50,\ .84,\ .57,\ .34,\ .64,\ .49] \frac{\mathrm{Ns}}{\mathrm{mm}}$. We approximate the characteristics connecting input pressure to actuation forces as linear. Its slope is estimated to be $180.682 \frac{\mathrm{Nmm}}{\mathrm{Bar}}$. See [4] for details on how these estimations are performed.

The soft robot's configurations are acquired at 100 Hz using a motion capture system (OptiTrack). Markers are placed on the robot to designate $\{S_i\}$. The measurements are mapped in q as in [20]. External forces are estimated using (14), with gain $\gamma = 0.5$. The model of uncertainty is trained using Algorithm 1. GPR training is done using `fitrgp` function from `MatLab2019b`'s `Statistics and Machine Learning Toolbox`. The dataset is comprised of input pressures and configurations while the arm is in the rest configuration and at steady-state after actuation by various inflation steps. First, the same actuator in all segments is inflated to 100 kPa. The procedure is repeated for all actuators (1 through 4). The same experiment is repeated, but with all the possible combinations of two neighboring actuators in each segment inflated to 100 kPa. The average norm of the prediction error is 1.27N without the GPR and 0.07N with it.

It is worth highlighting here that assessing the force estimation accuracy is beyond the scope of the present paper. Rather, we are interested in more qualitative measurements: detecting the presence and the persistence of a contact event, as well as the direction of the contact force, which is not easily achieved with current soft robotic sensors. We consider two contact scenarios. In Experiment 1, presented in Fig. 1, a table is placed in front of the robot and the robot actuates into its edge. Only Actuator 2 of each segment is inflated in Experiment 1. Two actuation pressures are considered: small (50 kPa) and medium (75 kPa). In the first case, no contact with the table is established, while in the second, the robot touches the table edge (see Fig. 1 b). Fig. 5(a) and (b) show the evolution of q for both conditions of Experiment 1. No major difference can be intuitively recognized between the two. Nonetheless, the algorithm is able to correctly detect and characterize the contact event. Figs. 5(c) and (d) show the contact forces $\hat{\tau}_{\mathrm{ext}}$ estimated by the algorithm. After a transient due to the influence of dynamic forces, the estimation converges to zero in the no contact case, while in the contact case, it is obviously non-zero. This can be translated in Cartesian forces through pre-multiplication or the pseudo-inverse of the transpose Jacobian. The result is shown in Fig. 5(e) and (f), where similar trends can be observed: a convergence of f_{ext} towards zero in the non-contact case, and non-zero f_{ext} in the contact case. Note that the force regressed during table contact points in the direction of $-y$, which is expected given the table's position and the direction of pushing (Fig. 1).

In Experiment 2, the soft robot is actuated by inflating the same two neighboring chambers in each segment and, while still swinging, pressed with a pole to simulate contact events (Fig. 6). The contact is applied twice by pushing the robot in the direction of the first quadrant of (x, y), and twice in the direction of the fourth quadrant (where quadrants are expressed in the global coordinates depicted in Fig. 6(i)). The experiment is repeated three times; in each one a different segment is touched with the pole (see Fig. 6(iii)-(viii)). This scenario is particularly challenging since the interaction is rapid and occurs while the robot is still into a transient due to the actuation. Nonetheless, the algorithm correctly detects the interactions any time they occur, with a 100% success rate, as shown in Figs. 7 and 8. Fig. 8(a,c,e) show the estimated torques. The peaks corresponds to interactions with the pole, while the small-amplitude, high-frequency transients, analogous to the ones observed in Experiment 1, are due to dynamic forces. Looking at these patterns, the contact events can be localized in time. Finally, an equivalent Cartesian representation can be derived, as for

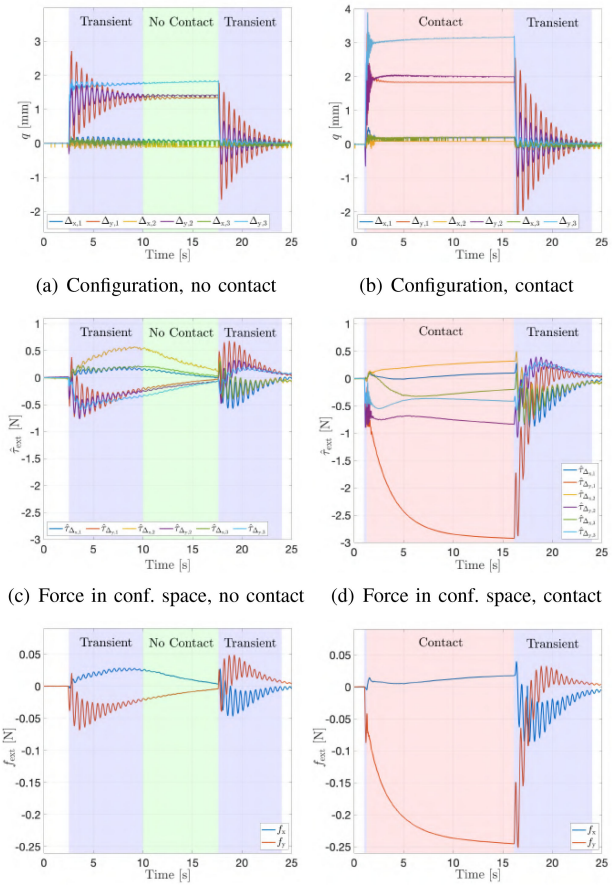

(a) Configuration, no contact (b) Configuration, contact

(c) Force in conf. space, no contact (d) Force in conf. space, contact

(e) Force in Cartesian, no contact (f) Force in Cartesian space, contact

Fig. 5. Results from Experiment 1. A table is placed in front of the robot. Panel (a) shows configuration when a small inflation is used for actuation and no contact occurs. Panel (b) shows the same when contact is established for high-pressure actuation. Panels (c,d) depict the estimated forces in the two conditions. In the no-contact case, $\hat{\tau}_{\text{ext}}$ converges to zero, while the result converges to large non-zero values during contact. Panels (e,f) show the estimated forces expressed in Cartesian coordinates.

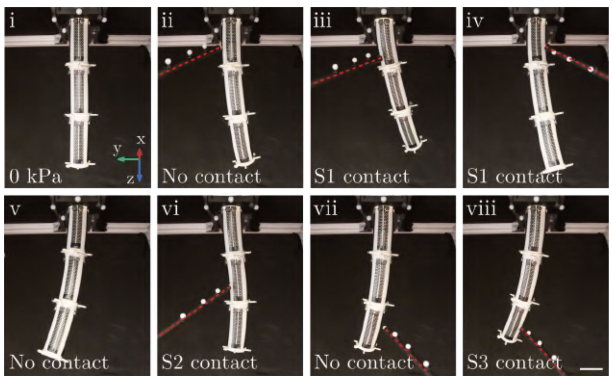

Fig. 6. Applying Contact Forces in Experiment 2. Stills from Experiment 2 are provided to show examples of the types of contact events we explored. This experiment begins with a soft robot at rest, as shown in (i). The robot is actuated to a pose like that in (ii) and perturbed or pressed with a pole (iii). The pole is highlighted for clarity with a red dashed line. We continue in this manner for various poses, as illustrated in the various examples provided in (iv) to (viii), which show the actuated robot before and after contact perturbations. Contact forces are applied to the first (iii, iv), second (vi), and third (viii) segments. (The reference frame provided in (i) is the same as that in Experiment 1, and the scale bar is 5 cm.)

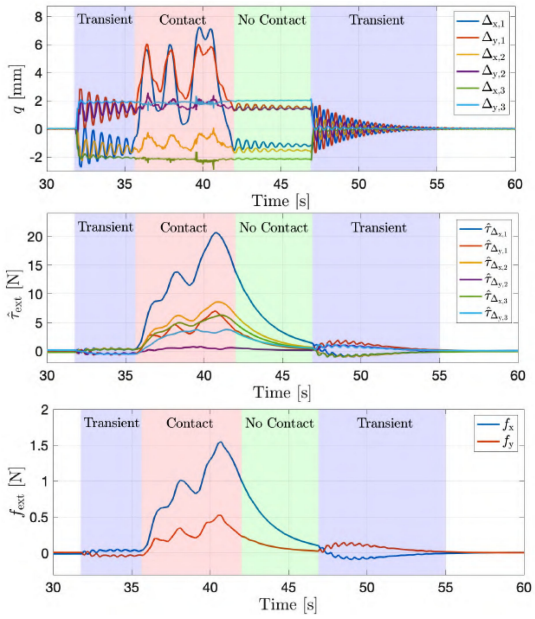

Fig. 7. Results from One Contact Event in Experiment 2. The robot's configuration (top) and estimated external forces expressed in configuration ($\hat{\tau}_{\text{ext}}$) (middle) and Cartesian (f_{ext}) coordinates (bottom) are provided for one actuation sequence. After the robot is actuated with a step inflation input at $t = 32$s, it is perturbed by pressing Segment 2 with a pole.

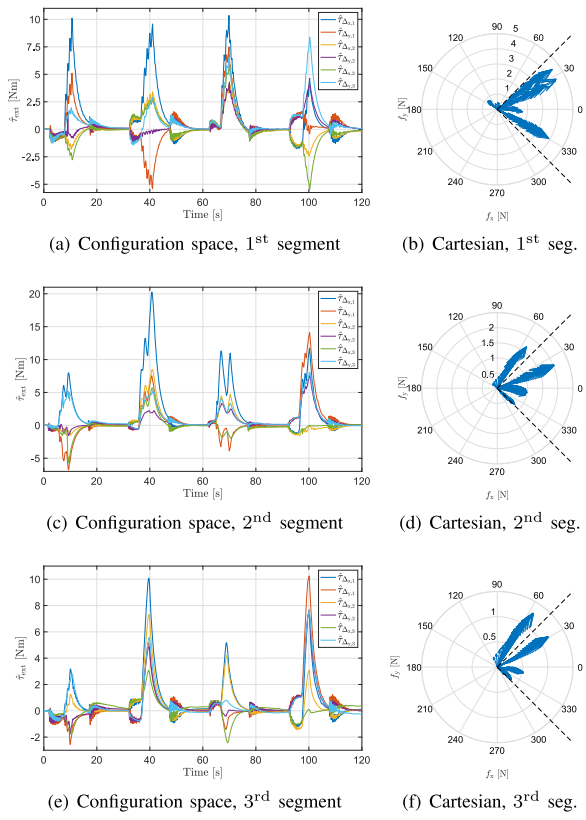

(a) Configuration space, 1st segment (b) Cartesian, 1st seg.

(c) Configuration space, 2nd segment (d) Cartesian, 2nd seg.

(e) Configuration space, 3rd segment (f) Cartesian, 3rd seg.

Fig. 8. Complete Results from Experiment 2. Panels (a,c,e) show the time evolution of $\hat{\tau}_{\text{ext}}$. These forces are mapped in Cartesian space in Panels (b,d,f). The data in Panels (a,b), (c,d), and (e,f) correspond to trials in which the robot is pressed on Segments 1, 2, and 3, respectively, as it is changing its pose. In each trial, the robot actuates from 0 kPa, is pressed on Segments 1/2/3, and deflated. Contacts happen twice each in the directions of $(+x, -y)$ and $(+x, +y)$ (of the global coordinates, see videos).

Experiment 1 (see Fig. 8(b,d,f)), where the four contact events discussed above are distinctly shown. Two set of arrows are expected to point in the direction of the first quadrant and two in the direction of the fourth. This behavior is clearly observed in Fig. 8(b,d,f). To give quantitative measurements of this effect, we evaluated the normalized scalar product of the first set of vectors with $\frac{1}{\sqrt{2}}[1, \quad 1]^{\mathrm{T}}$. Note that vectors of norm smaller than 0.05N have been discarded in this analysis. Additionally, data extracted from another complete set of analogous experiments are added to the dataset, to increase the statistical significance of the analysis. The average alignment evaluated with this method is 0.95 (where 1 is perfectly aligned, and 0 is orthogonal), with standard deviation 0.12. Another analysis comparing the second set of vectors with $\frac{1}{\sqrt{2}}[1, \quad -1]^{\mathrm{T}}$ results in an average alignment of 0.92 with standard deviation 0.08. These results confirm that our algorithm enables estimation of the direction of contact forces applied to the soft robot.

V. Conclusion and Discussion

The methods presented here address a key need in soft robotic control: a means for estimating contact forces acting on a soft robot body during environmental interactions. Our approach combines the sample efficiency, interpretability, and analytical nature afforded by our new class of model-based observers, with the opportunities machine learning uniquely enables in accounting for the various uncertainties and non-modeled parameters soft robots encounter in reality. Again, our methods are only valid for low frequency contact events and are not appropriate for realizing soft robots that can quickly react to rapid, unexpected interactions. Achieving this goal would require the full use of all the dynamic terms in (1), as well as to measure $\dot{q}$. Lastly, while our algorithm currently uses posture information determined from motion capture for estimating external forces, we are now working to leverage system-integrated soft proprioceptive sensors, as the ones shown in Figs. 1 and 6. Still, our current algorithms can accompany soft sensorized robots with proprioceptive and tactile sensing capabilities [25] to obtain greater detail about interactions. We are now exploring other machine learning techniques as alternatives to GPRs, such as LSTM networks [24], and Support Vector Regressors [26]. We also plan to improve the ability of the nominal disturbance observer to deal with dynamic conditions by considering polynomial curvature models [27]. Our ultimate goal is to use these methods for operating soft robots involved in continuous, controlled interactions with unstructured environments. We are particularly interested in enabling soft robots designed for surgical, exploratory, and collaborative applications to better perceive and act on their interactions with the outside world.

Acknowledgment

The authors would like to thank Dr. Robert Katzschmann for original designs of the 3D soft robotic arm.

References

[1] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.

[2] E. Coevoet, A. Escande, and C. Duriez, "Optimization-based inverse model of soft robots with contact handling," *IEEE Robot. Autom. Lett.*, vol. 2, no. 3, pp. 1413–1419, Jul. 2017.

[3] T. G. Thuruthel, Y. Ansari, E. Falotico, and C. Laschi, "Control strategies for soft robotic manipulators: A survey," *Soft Robot.*, vol. 5, no. 2, pp. 149–163, 2018.

[4] C. Della Santina, R. K. Katzschmann, A. Bicchi, and D. Rus, "Model-based dynamic feedback control of a planar soft robot: Trajectory tracking and interaction with the environment," *Int. J. Robot. Res.*, vol. 39, no. 4, pp. 490–513, 2020.

[5] H. Wang, M. Totaro, and L. Beccai, "Toward perceptive soft robots: Progress and challenges," *Adv. Sci.*, vol. 5, no. 9, 2018, Art. no. 1800541.

[6] A. Bicchi, "Intrinsic contact sensing for soft fingers," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1990, pp. 968–973.

[7] M. T. Gillespie, C. M. Best, and M. D. Killpack, "Simultaneous position and stiffness control for an inflatable soft robot," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2016, pp. 1095–1101.

[8] A. Ataka *et al.*, "Real-time pose estimation and obstacle avoidance for multi-segment continuum manipulator in dynamic environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 2827–2832.

[9] D. C. Rucker and R. J. Webster, "Deflection-based force sensing for continuum robots: A probabilistic approach," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2011, pp. 3764–3769.

[10] Y. Chen, L. Wang, K. Galloway, I. Godage, N. Simaan, and E. Barth, "Modal-based kinematics and contact detection of soft robots," *Soft Robot.*, 2020.

[11] T. G. Thuruthel, B. Shih, C. Laschi, and M. T. Tolley, "Soft robot perception using embedded soft sensors and recurrent neural networks," *Sci. Robot.*, vol. 4, no. 26, 2019, Art. no. eaav1488.

[12] E. Magrini, F. Flacco, and A. De Luca, "Estimation of contact forces using a virtual force sensor," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 2126–2133.

[13] S. Haddadin, A. De Luca, and A. Albu-Schäffer, "Robot collisions: A survey on detection, isolation, and identification," *IEEE Trans. Robot.*, vol. 33, no. 6, pp. 1292–1312, Dec. 2017.

[14] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Mach. Learn.*. Berlin, Germany: Springer-Verlag, 2003, pp. 63–71.

[15] R. K. Katzschmann, C. Della Santina, Y. Toshimitsu, A. Bicchi, and D. Rus, "Dynamic motion control of multi-segment soft robots using piecewise constant curvature matched with an augmented rigid body model," in *Proc. 2nd IEEE Int. Conf. Soft Robot.*, 2019, pp. 454–461.

[16] B. Akgun, M. Cakmak, J. W. Yoo, and A. L. Thomaz, "Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective," in *Proc. 7th Annu. ACM/IEEE Int. Conf. Human-Robot Interact.*, 2012, pp. 391–398.

[17] K. Hsiao, L. Kaelbling, and T. Lozano-Pérez, "Task-driven tactile exploration," *Robot.: Sci. Syst. VI*, pp. 225–233, 2010.

[18] S. M. H. Sadati *et al.*, "Stiffness imaging with a continuum appendage: Real-time shape and tip force estimation from base load readings," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2824–2831, Apr. 2020.

[19] R. J. Webster III and B. A. Jones, "Design and kinematic modeling of constant curvature continuum robots: A review," *Int. J. Robot. Res.*, vol. 29, no. 13, pp. 1661–1683, 2010.

[20] C. Della Santina, A. Bicchi, and D. Rus, "On an improved state parametrization for soft robots with piecewise constant curvature and its use in model based control," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 1001–1008, Apr. 2020.

[21] V. F. Zaitsev and A. D. Polyanin, *Handbook of Exact Solutions for Ordinary Differential Equations*. London, U.K.: Chapman & Hall, 2002.

[22] C. Della Santina, L. Pallottino, D. Rus, and A. Bicchi, "Exact task execution in highly under-actuated soft limbs: An operational space based approach," vol. 4, no. 3, pp. 2508–2515, 2019.

[23] B. Oksendal, *Stochastic Differential Equations: An Introduction With Applications*. Berlin, Germany: Springer, 2013.

[24] R. L. Truby, C. Della Santina, and D. Rus, "Distributed proprioception of 3D configuration in soft, sensorized robots via deep learning," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 3299–3306, Apr. 2020.

[25] R. L. Truby *et al.*, "Soft somatosensitive actuators via embedded 3D printing," *Adv. Mater.*, vol. 30, no. 15, 2018, Art. no. 1706383.

[26] A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statist. Comput.*, vol. 14, no. 3, pp. 199–222, 2004.

[27] C. Della Santina and D. Rus, "Control oriented modeling of soft robots: The polynomial curvature case," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 290–298, Apr. 2020.