

A Soft Robotic System Automatically Learns Precise Agile Motions Without Model Information

Simon Bachhuber^{1†}, Alexander Pawluchin^{2†}, Arka Pal³, Ivo Boblan², Thomas Seel⁴

Abstract—Many application domains, e.g., in medicine and manufacturing, can greatly benefit from pneumatic Soft Robots (SRs). However, the accurate control of SRs has remained a significant challenge to date, mainly due to their nonlinear dynamics and viscoelastic material properties. Conventional control design methods often rely on either complex system modeling or time-intensive manual tuning, both of which require significant amounts of human expertise and thus limit their practicality. In recent works, the data-driven method, Automatic Neural ODE Control (ANODEC) has been successfully used to – fully automatically and utilizing only input-output data – design controllers for various nonlinear systems *in silico*, and without requiring prior model knowledge or extensive manual tuning. In this work, we successfully apply ANODEC to automatically learn to perform agile, non-repetitive reference tracking motion tasks in a real-world SR and within a finite time horizon. To the best of the authors’ knowledge, ANODEC achieves, for the first time, performant control of a SR with hysteresis effects from only 30s of input-output data and without any prior model knowledge. We show that for multiple, qualitatively different and even out-of-training-distribution reference signals, a single feedback controller designed by ANODEC outperforms a manually tuned PID baseline consistently. Overall, this contribution not only further strengthens the validity of ANODEC, but it marks an important step towards more practical, easy-to-use SRs that can automatically learn to perform agile motions from minimal experimental interaction time.

I. INTRODUCTION

Soft Robots (SRs), including Pneumatic Soft Actuators (PSAs) and Pneumatic Artificial Muscles (PAMs), are gaining significant interest in diverse (bio)medical and industry application domains [1], [2]. Their rising popularity can be attributed to their inherent soft characteristics, derived from the use of flexible and deformable materials, which enables SRs to offer unique capabilities and adaptability in uncertain environments [3]. These attributes prove particularly advantageous in scenarios such as rehabilitation support [4] and human-robot interaction [5], where providing a safe interaction environment without external sensors and safety mechanisms is crucial.

Unfortunately, the very features that give PSAs their unique advantages also introduce several challenges for their accurate control, a dilemma that is yet to be fully addressed

[†]Both S.B. and A.P. contributed equally to the manuscript.

¹S.B. is part of the Department Artificial Intelligence in Biomedical Engineering, FAU Erlangen-Nürnberg, 91052 Erlangen, Germany

²A.P. and I.B. are part of the CoRoLab Berlin, University of Applied Sciences and Technology, 13353 Berlin, Germany

³Ar.P. is part of the School of Mechanical Sciences, Indian Institute of Technology Bhubaneswar, 752050 Odisha, India

⁴T.S. is part of the Institute of Mechatronic Systems, Leibniz Universität Hannover, 30167 Hannover, Germany

Corresponding author: S.B. (simon.bachhuber@fau.de)

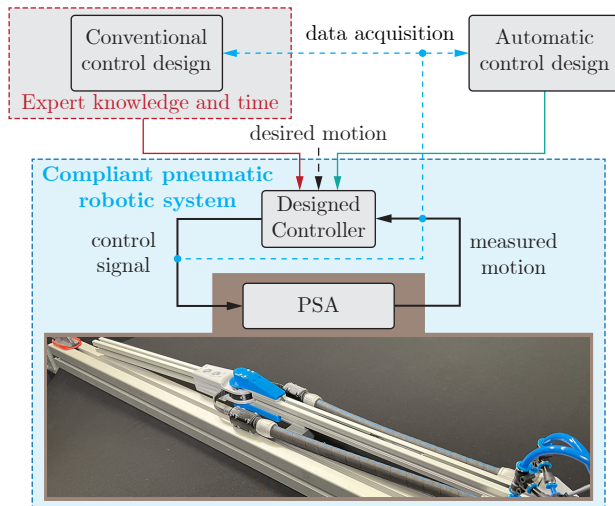


Fig. 1. Comparison between conventional controller design, reliant on human expertise, and an automatic, data-driven control structure for Pneumatic Soft Actuators (PSAs). Both controller design paradigms observe the PSA’s control signal and resulting measured motion, and subsequently design position feedback controllers that enable the measured motion of the PSA to track the desired motion. In this work we show that the data-driven method, Automatic Neural ODE Control (ANODEC), enables automatic control design for a PSA.

[3]. These challenges are for example: a) strong nonlinearities as the dynamics of PSAs are tightly coupled to the actuator’s velocity which leads to lower control accuracy and overshoot, and b) creeping and hysteresis (cf. Figure 3, bottom subplot) due to the viscoelasticity and time-dependent properties of the soft material.

Despite these challenges, several methods have been previously applied for the control of PSAs. Model-based methods have been applied for PSAs such as, e.g., [6]–[9] but they typically require a detailed system model for accurate control and therefore involve extensive system modeling combined with human expertise.

Data-driven methods can alleviate the need for extensive system modeling and have been applied for PSAs. In [10], [11], Koopman operators are used to achieve reference tracking of a SR, however, they make restrictive assumptions such as either weak hysteresis characteristics or non-agile reference trajectory targets. In [12]–[14], a DNN-based MPC was employed to control SRs but they typically focus on non-agile motions and within a computationally intensive framework. In [15], an ODE-based kinematic model is learned and combined with MPC in order to learn to perform non-agile, quasistatic motions. In [16]–[18], Iterative Learning Control

(ILC) in combination with different feedback controllers are used for the control of PSAs. However, ILC is tailored to repetitive motions and requires a series of iterations. Only in [18], deep learning is used to interpolate from previously learned trajectories and to achieve a broader range of motions, but they require large amounts of system interaction time to learn a single motion.

Recently, the data-driven method, Automatic Neural ODE Control (ANODEC) has been shown to – fully automatically – design feedback controllers for various nonlinear systems *in silico*, and without requiring prior model knowledge or extensive manual tuning. Unlike various other data-driven methods, e.g., [19]–[24], ANODEC requires neither the full state to be observed nor does it require an approximate system model.

In this contribution, we successfully apply ANODEC to automatically learn to perform agile, non-repetitive reference tracking motion tasks in a real-world PSA within a finite time horizon. To the best of the authors’ knowledge, ANODEC achieves, for the first time, performant control of a PSA with hysteresis effects (cf. Figure 3) from only 30 s of input-output data and without any prior model knowledge. Additionally, the single feedback controller designed by ANODEC enables tracking of arbitrary reference signals and is validated on multiple, qualitatively different, and even out-of-training-distribution reference signals. Moreover, it is shown that ANODEC outperforms a manually tuned PID baseline, which similar to ANODEC, requires no prior model knowledge and does not require the system state to be observed.

Overall, it is shown that ANODEC allows for automatic control design (cf. Figure 1) and can enable PSAs to learn to perform agile motions from only a parsimonious amount of experimental interaction time and without requiring any prior model knowledge.

II. PNEUMATIC SOFT ACTUATOR

We consider a PSA, as shown in Fig. 2, that typically consists of a rigid or continuous kinematic structure and is driven by an antagonistic pair of PAMs. The muscles are attached via a belt and form a 1-Degree-of-Freedom (DOF) arm. The hinge joint angle represents the system output $\varphi \in \mathbb{R}$ and it is limited by the geometrical configuration of the muscle contraction and diameter of the pulley such that the range of motion $\varphi \in [\varphi_{\min}, \varphi_{\max}]$ where $\varphi_{\min/\max} \in \mathbb{R}$.

Each of the PAMs exhibits undesired parasitic hysteresis and creep properties that strongly influence the nonlinear characteristic force map $F_{1,2} \in \mathbb{R}$. Given the unknown forces F_1 and F_2 , we instead control the pressure in the bellows. For this purpose, we use two model-based pressure controllers, which are designed to follow the given desired pressure $p_{d,1} \in \mathbb{R}$ and $p_{d,2} \in \mathbb{R}$.

Based on this system description, the system output is given by

$$\varphi(t) = \tilde{\Psi}(p_{d,1}(t' < t), p_{d,2}(t' < t)) \quad \forall t \in [0, T], \quad (1)$$

where $\tilde{\Psi}$ captures the system’s dynamics (including the two pressure controllers), $T \in \mathbb{R}$ is the finite trial duration and




Fig. 2. The PSA including two PAMs is controlled using two pressure controllers, which implement a SISO control strategy through the mean and difference pressure approach. As a result, the single control input u controls both desired difference pressures $p_{d,1/2}$, and the single system output is the hinge joint angle φ .

where $p_{d,1/2}(t' < t)$ is used to denote that $\varphi(t)$ depends on all desired pressures up to time t , i.e., $p_{d,1/2}(t') \forall t' < t$.

Since the PSA has two pressure inputs, we couple them using the mean pressure approach in (2) to reduce the input dimension which then yields a Single-Input Single-Output (SISO) system. This is done by defining a mean pressure $p_m = \text{const.} \in \mathbb{R}$ and a difference pressure $\Delta p_d = p_{d,1} - p_{d,2}$, i.e.,

$$p_{d,1} = p_m + 0.5u \quad \text{and} \quad p_{d,2} = p_m - 0.5u. \quad (2)$$

The difference pressure becomes the system’s single control input $u \in \mathbb{R}$, i.e., $u := \Delta p_d$, which is limited by the maximum physical permissible pressure range $u_{\min} \in \mathbb{R}$ and $u_{\max} \in \mathbb{R}$. Finally, using (1) and (2), the system in SISO configuration is given by

$$\varphi(t) = \Psi(u(t' < t)) \quad \forall t \in [0, T]. \quad (3)$$

where the functional Ψ is called the system input-output map.

In this work, we aim to design a feedback controller that manipulates $u(t)$ to let $\varphi(t)$ follow a given time-varying reference signal $\varphi_d(t) \in \mathbb{R}$. Thus, we seek to find a controller Φ that maps the desired angle $\varphi_d(t)$ and the measured output $\varphi(t)$ to the input vector $u(t)$, that is,

$$u(t) = \Phi(\varphi_d(t' < t), \varphi(t' < t)) \quad \forall t \in [0, T], \quad (4)$$

such that it minimizes the tracking error between the output and the reference signal, i.e.,

$$\Phi^* = \arg \min_{\Phi} \int_0^T \|\varphi_d(t) - \varphi(t)\|_2 dt \quad (5)$$

over the finite trial duration T . Note that the desired feedback controller measures the system output and not the state and that an arbitrary reference signal can be provided in real time and must not be provided ahead of time. However, we assume that the system captured by Ψ is repeatedly reset to some trial-invariant initial state before each trial, and that the open-loop system can be safely probed for input-output data. Note that the latter assumption is not only true for open-loop stable systems but can also be true for unstable system where

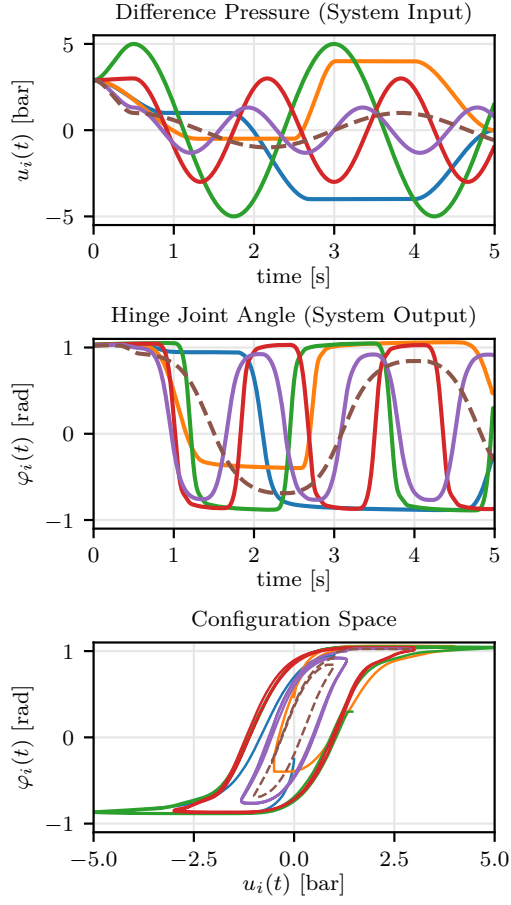


Fig. 3. Training data for ANODEC consists of five input-output pairs, each of a five seconds length, that are gathered from the experimental PSA. One additional input-output pair (dashed line) is collected and used as validation data and to prevent model overfitting. The feasible input and output intervals of the experimental PSA are $u(t) \in [-6, 6]$ bar and $\varphi(t) \in [-1, 1]$ rad, respectively.

the instability acts on large timescales relative to the finite trial duration, i.e., it is sufficiently slow.

III. AUTOMATIC NEURAL ODE CONTROL (ANODEC)

Internally, the ANODEC algorithm [25] designs controllers in a two-step approach of model learning from input-output data pairs, and controller learning that utilizes the learned model.

A. Data Requirements

To apply ANODEC, a dataset \mathcal{D} that comprises of six pairs of input-output data over time, each corresponding to one trial of length $T = 5$ s, is gathered from the system input-output map Ψ . The data is recorded at 100 Hz. Fig. 3 shows the six input-output data pairs.

The model learning (cf. Section III-B) then uses the dataset

$$\mathcal{D} = \{(u_i(t), \varphi_i(t)) | i \in 1 \dots 6\}. \quad (6)$$

The sixth pair of input-output data is split from the dataset \mathcal{D} and used as validation data. The validation data is used to ensure that the learned model is accurate and that the

amount of collected data is sufficient for an accurate model. Thus, in total ANODEC requires 30 s of interaction time with the PSA. Input trajectories $u(t)$ for gathering the data and thus probing the system are drawn from a sinusoidal and a spline function generator. These functions are provided as pseudocode in Algorithm 1. Before each trial starts, the system enters a saturation state such that $\varphi(t = 0) \approx \varphi_{\max}$. This procedure is required as otherwise there might be state ambiguity due to the hysteresis of the PAMs. Note that the system probing and model learning step (cf. Section III-B) can easily be extended to an iterative process that switches between system probing and model learning until a sufficiently accurate model is obtained. This way ANODEC can be extended to a fully autonomous control design algorithm that iteratively improves the performance of the designed feedback controller.

B. Model Learning

ANODEC approximates Ψ in (3) by a neural ODE that is learned from pairs of experimental input-output trajectories $u_i(t), \varphi_i(t)$ (cf. Section III-A). The neural ODE that approximates Ψ is given by

$$\begin{aligned} \frac{d\xi^{(m)}(t)}{dt} &= \tanh\left(A_1^{(m)}\left(\xi^{(m)}\tau(t), u(t)\right)^\top + b_1^{(m)}\right), \quad (7) \\ \hat{\varphi}(t) &= A_2^{(m)}\xi^{(m)}(t) + b_2^{(m)}, \end{aligned}$$

where the latent state vector $\xi^{(m)} \in \mathbb{R}^9$, and with model parameters $A_1^{(m)} \in \mathbb{R}^{9 \times 10}$, $b_1^{(m)} \in \mathbb{R}^9$, $A_2^{(m)} \in \mathbb{R}^{1 \times 9}$, $b_2^{(m)} \in \mathbb{R}$. We denote the vector of all model parameters by $\theta^{(m)} \in \mathbb{R}^{109}$ and use supervised learning to estimate and then optimize for these parameters

$$\theta^{(m)*} = \arg \min_{\theta^{(m)}} \sum_{i=1}^5 \int_0^T \|\varphi_i(t) - \hat{\varphi}_i(t)\|_2 dt. \quad (8)$$

C. Controller Learning

As a second step, ANODEC designs a feedback controller using the trained model with frozen parameters $\theta^{(m)*}$. Internally, it performs a very large number of forward simulations of the closed-loop dynamics (of trained model and feedback controller) with randomly drawn step reference signals.

ANODEC represents the controller Φ as a neural ODE, given by

$$\begin{aligned} \frac{d\xi^{(c)}(t)}{dt} &= A_1^{(c)}\left(\xi^{(c)}\tau(t), \varphi(t), \varphi_d(t)\right)^\top + b_1^{(c)}, \quad (9) \\ \bar{u}(t) &= \tanh\left(A_2^{(c)}\xi^{(c)}(t) + b_2^{(c)}\right), \\ u(t) &= (u_{\max} - u_{\min})(\bar{u}(t) \cdot 0.5 + 0.5) + u_{\min}, \end{aligned}$$

where the latent state vector $\xi^{(c)} \in \mathbb{R}^5$, and with controller parameters $A_1^{(c)} \in \mathbb{R}^{5 \times 7}$, $b_1^{(c)} \in \mathbb{R}^5$, $A_2^{(c)} \in \mathbb{R}^{1 \times 5}$, $b_2^{(c)} \in \mathbb{R}$. We denote the vector of controller parameters by $\theta^{(c)} \in \mathbb{R}^{46}$.

ANODEC then closes the loop between the neural ODEs that approximate the system input-output map Ψ (7) and the controller Φ (9). The resulting closed-loop ODE takes




Fig. 4. Two experimental setups of a pneumatic arm with a single DOF. In both setups, two PAMs (pneumatic artificial muscles, black tubes) are used as an antagonistic pair to control the arm's forces and position. The upper setup shows the simplest configuration without the influence of gravity and external load. The lower setup is loaded with an external weight of 0.6 kg, with a lever arm of 0.25 m oriented against gravity. Both arms are fixed to the ground to prevent undesired movements.

the reference $\varphi_d(t)$ as an input and outputs $\hat{\varphi}(t)$, and we optimize the controller parameters $\theta^{(c)}$ via

$$\theta^{(c)*} = \arg \min_{\theta^{(c)}} \left(\lambda^{(c)} \|\theta^{(c)}\|_2 + \mathbb{E}_{\varphi_d \sim \mathcal{U}(\varphi_{\min}, \varphi_{\max})} \left[\int_0^T \|\varphi_d - \hat{\varphi}(t)\|_2 dt \right] \right) \quad (10)$$

where $\lambda^{(c)} = 4 \cdot 10^{-4}$ is a small regularization weight, and the reference φ_d is a constant value drawn uniformly from $[\varphi_{\min}, \varphi_{\max}]$.

D. Time Integration & Optimization

We use Runge-Kutta of fourth order (RK4) to obtain a numerical solution of the time integration in (8) and (10), together with the initial condition $\xi^{(m)}(t=0) = 0$ and $\xi^{(m)}(t=0) = \xi^{(c)}(t=0) = 0$, respectively.

For model training, the five input-output pairs from the training dataset \mathcal{D} are combined into a single model training batch. For controller training, the expectation operator $\mathbb{E}_{\varphi_d \sim \mathcal{U}(\varphi_{\min}, \varphi_{\max})}$ is estimated by randomly drawing 50 constant reference values at every optimizer step.

Both (8) and (10) are optimized using the Adam optimizer with a learning rate of 1×10^{-3} and global gradient L2 norm clipping of 1.0. Gradients are computed using backpropagation. The model is trained for 50 000 optimizer steps and the single validation input-output pair is used for early stopping and thus prevent overfitting during the model learning. The controller is trained for 12 000 optimizer steps.

IV. EXPERIMENTAL VALIDATION

In this section, we validate ANODEC (cf. Section III) on an experimental PSA (cf. Section II). To showcase the generalization capabilities of ANODEC in Section IV-C, we

propose unseen reference signal distributions that, in Section IV-D, will be used to validate ANODEC's performance compared to a PID controller baseline (cf. Section IV-B). The corresponding demonstration video to the presented results in Fig. 5 can be found [here](#).

A. Experimental setup

We evaluate ANODEC's performance on a pneumatic compliant robotic arm which should solve a pick and place motion task on a given reference trajectory. The arm, as depicted in Fig. 4, consists of both an upper and a lower arm. The upper arm carries the PAMs responsible for controlling the motion of the lower arm. The hinge joint has a range of motion within $\varphi \in [-1, 1]$ rad and connects both PAMs (DMSP-10-250N) via a constant pulley. The angle is measured using a MLX90316 rotary position sensor. The upper setup (Setup 1) illustrates the configuration without external influences such as gravity and load. The lower setup (Setup 2) shows the arm loaded with a 0.6 kg weight on a 0.25 m lever arm against gravity. Due to the non-negligible mass of the lower arm and the highly agile target motions, the upper arm is securely fixed to the ground to prevent undesired movements.

To generate forces, the PAMs are driven using a model-based pressure controllers and two proportional directional control valves MPYE-5-M5-010-B for controlling the mass flow. The controlled pressure is measured using two SPTE-P10R-S4-B-2.5K pressure transmitters. The supply pressure was set to 8 bar. This results in a feasible system input range $u \in [-6, 6]$ bar.

The pressure controllers runs on an embedded computer with a sampling rate of 200 Hz. This system is connected through Ethernet and ROS to a computer (Intel i7 7700) that controls both ANODEC and PID approaches in Python with a control rate of 100 Hz.

B. Baseline Controller

As a baseline, we compare the performance of ANODEC to a manually tuned PID controller, which similar to ANODEC, requires no prior model knowledge and does not require the system state to be observed. Eleven trials were required (equaling 55 s of interaction time) for manual tuning and to achieve a sufficiently performant PID controller. Since we only control the pressure and have no assumptions regarding the force-pressure relation, an integral part is needed to compensate for the PAMs force nonlinearity and hysteresis.

The tuned gain parameters are $k_p = 2$, $k_i = 30$ and since the derivative term of the PID controller is highly sensitive to the sensor noise from the angle sensor, leading rapidly to undesired instabilities, the D term is set to zero. The PID controller's output u was clipped to stay within the feasible system input range of $[u_{\min}, u_{\max}]$. This is not required for ANODEC since its output remains within the feasible input range by design, see (9).

C. Generalization to Tracking of Unseen Reference Signals

A common criticism of data-driven or machine-learning-based solutions is that they might only perform well in




Fig. 5. Performance comparison of the automatic ANODEC and a manually tuned PID controller baseline in Setup 1 for one exemplary reference signals drawn from the three reference signal distributions: Step, double step, and smooth. Even on reference signals that were not present during training (double step, smooth), ANODEC is able to consistently outperform the PID baseline and achieves a lower RMSE tracking error while requiring less experimental interaction time. Video (<https://youtu.be/7HkXKy0WuRw>) that showcases these trials.

TABLE I

REFERENCE TRACKING RMSE (IN DEGREES) OF PID AND ANODEC FOR THREE REFERENCE SIGNAL DISTRIBUTIONS AND TWO SYSTEMS: STEPS, DOUBLE STEPS, AND CUBIC SPLINES. FOR EACH DISTRIBUTION, N DISTINCT REFERENCE SIGNALS ARE DRAWN AND USED TO ESTIMATE THE MEAN-RMSE AND ONE STANDARD DEVIATION.

REFERENCES	PID	ANODEC
SETUP 1		
STEPS ($N = 2$)	12.89 ± 6.14	10.03 ± 4.93
DOUBLE STEPS ($N = 2$)	12.86 ± 2.66	11.08 ± 1.37
CUBIC SPLINES ($N = 12$)	10.01 ± 1.44	4.48 ± 1.11
SETUP 2		
STEPS ($N = 2$)	4.00 ± 0.45	5.54 ± 0.53
DOUBLE STEPS ($N = 2$)	9.50 ± 4.03	6.81 ± 0.75
CUBIC SPLINES ($N = 4$)	4.56 ± 0.81	5.01 ± 0.79

the cases covered by the training dataset. To verify that ANODEC generalizes beyond the step reference signals (used internally by ANODEC, see (10)), we test the performance on two additional, qualitatively different reference signal distributions: Double-step and cubic-spline reference trajectories. In total, 16 different reference signals are used to validate ANODEC. One exemplary reference signal for each of the three reference signal distributions is illustrated in Fig. 5 (dashed black line).

D. ANODEC Reference Tracking Performance

We validate the proposed method (cf. Section III) and compare to the baseline controller (cf. Section IV-B) for 16 different reference signals (cf. Section IV-C) on the experimental setup of the PSA (cf. Section IV-A).

ANODEC is able to outperform the PID controller baseline for all three reference signal distributions, as can be seen in Table I. Overall, ANODEC consistently outperforms the PID controller baseline in all 16 reference signals individually, and achieves a, on average, $\approx 45.9\%$ lower RMSE tracking error across all reference signals while requiring less experimental interaction time (30 s versus 55 s). For each reference signal distribution, one exemplary reference signal

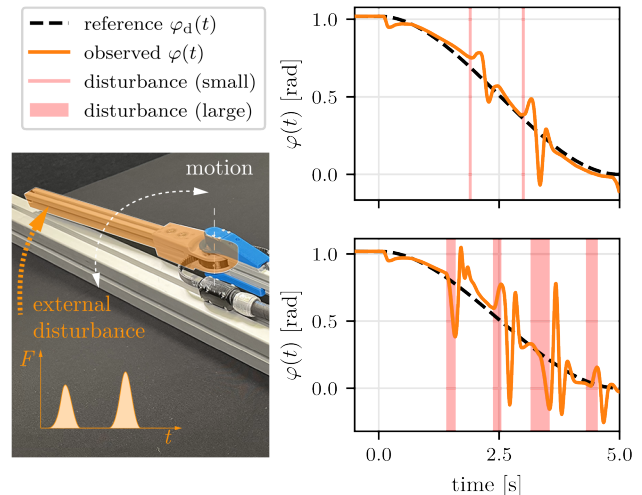


Fig. 6. Tracking performance of ANODEC in two disturbed trials. In the first trial (top) there are two small disturbances (poking the endeffector with a stick) whereas in the second trial (bottom) there are four large disturbances (grabbing and briefly holding the endeffector). ANODEC remains stable and behaves as expected in both trials.

is drawn. The performance of ANODEC compared to the PID controller baseline is shown in Figure 5.

Additionally, in Figure 6 we show that ANODEC is stable even if the trials are heavily disturbed.

These results show that ANODEC enables automatic yet competitive controller design for a PSA which, in contrast to conventional methods, requires no human expertise.

E. System Variation

To demonstrate the applicability of ANODEC to a second experimental system (Setup 2), we modify the experimental setup as it is described in Section IV-A. This variation significantly alters the system dynamics because the arm now rotates vertically, and the attached mass introduces a strong nonlinear term due to gravity.

The trial duration is extended from 5 s to 8 s and, as before, we record six input-output pairs (five for training, one for validation), now resulting in 48 seconds of data. For this second use case, ANODEC successfully designs a

competitive controller. The achieved RMSEs are summarized in Table I.

V. CONCLUSION

In this contribution, we have demonstrated that the data-driven method ANODEC can enable PSAs to – fully automatically – learn to perform agile, non-repetitive motions, from only 30s of experimental interaction time. Utilizing only input-output data and without any prior model knowledge, ANODEC designs a single feedback controller that enables the tracking of arbitrary reference signals, which is then validated on multiple, qualitatively different, and even out-of-training-distribution reference signals. Specifically, it is shown that ANODEC can outperform a manually tuned PID baseline and achieves an up to 45.9% lower RMSE tracking error. Overall, this contribution not only reinforces the validity of ANODEC but also represents a significant advancement towards developing practical, user-friendly PSAs capable of learning to perform agile movements with only a minimal amount of experimental interaction time.

Future work will involve validation on different experimental PSAs, including PSAs with multiple DOFs, and increasing the autonomy of ANODEC by, e.g., self-tuning of involved hyperparameters (such as $\lambda^{(c)}$) and autonomous data acquisition.

REFERENCES

- [1] J. Walker, T. Zidek, C. Harbel, *et al.*, “Soft robotics: A review of recent developments of pneumatic soft actuators,” *Actuators*, vol. 9, no. 1, 2020. DOI: 10.3390/act9010003.
- [2] T.-L. Habich, J. Haack, M. Belhadj, D. Lehmann, T. Seel, and M. Schappler, “Sponge: Open-source designs of modular articulated soft robots,” *IEEE Robotics and Automation Letters*, vol. 9, no. 6, 2024. DOI: 10.1109/LRA.2024.3388855.
- [3] P. Polygerinos, N. Correll, S. A. Morin, *et al.*, “Soft robotics: Review of fluid-driven intrinsically soft devices; manufacturing, sensing, control, and applications in human-robot interaction,” *Advanced Engineering Materials*, vol. 19, no. 12, 2017. DOI: <https://doi.org/10.1002/adem.201700016>.
- [4] M. Martens, J. Zawatzki, T. Seel, and I. Boblan, “A pneumatic-muscle-actuator-driven knee rehabilitation device for cam therapy,” in *EMBC 2019*, 2019. DOI: 10.1109/EMBC.2019.8856526.
- [5] D. Müller, A. Raisch, A. Hildebrandt, and O. Sawodny, “Nonlinear model based dynamic control of pneumatic driven quasi continuum manipulators,” in *2020 IEEE/SICE SII*, 2020. DOI: 10.1109/SII46433.2020.9026206.
- [6] C. Della Santina, C. Duriez, and D. Rus, “Model-based control of soft robots: A survey of the state of the art and open challenges,” *IEEE Control Systems Magazine*, vol. 43, no. 3, 2023. DOI: 10.1109/MCS.2023.3253419.
- [7] T. George Thuruthel, Y. Ansari, E. Falotico, and C. Laschi, “Control strategies for soft robotic manipulators: A survey,” *Soft Robotics*, vol. 5, 2018. DOI: 10.1089/soro.2017.0007.
- [8] V. Falkenhahn, *Modellierung und modellbasierte Regelung von Kontinuum-Manipulatoren* (Berichte aus dem Institut für Systemdynamik Universität Stuttgart). Aachen: Shaker Verlag, 2017, vol. 32.
- [9] M. Martens, T. Seel, J. Zawatzki, and I. Boblan, “A novel framework for a systematic integration of pneumatic-muscle-actuator-driven joints into robotic systems via a torque control interface,” *Actuators*, vol. 7, no. 4, 2018. DOI: 10.3390/act7040082.
- [10] D. A. Haggerty, M. J. Banks, E. Kamenar, *et al.*, “Control of soft robots with inertial dynamics,” *Science Robotics*, vol. 8, no. 81, 2023. DOI: 10.1126/scirobotics.add6864.
- [11] D. Bruder, X. Fu, R. B. Gillespie, C. D. Remy, and R. Vasudevan, “Data-driven control of soft robots using koopman operator theory,” *IEEE Transactions on Robotics*, vol. 37, no. 3, 2021. DOI: 10.1109/TRO.2020.3038693.

- [12] M. T. Gillespie, C. M. Best, E. C. Townsend, D. Wingate, and M. D. Killpack, “Learning nonlinear dynamic models of soft robots for model predictive control with neural networks,” in *Robosoft 2018*, 2018. DOI: 10.1109/ROBOSOFT.2018.8404894.
- [13] C. Johnson, T. Quackenbush, T. Sorensen, D. Wingate, and M. Killpack, “Using first principles for deep learning and model-based control of soft robots,” *Frontiers in Robotics and AI*, vol. 8, 2021. DOI: 10.3389/frobt.2021.654398.
- [14] P. Hyatt, D. Wingate, and M. Killpack, “Model-based control of soft actuators using learned non-linear discrete-time models,” *Frontiers in Robotics and AI*, vol. 6, 2019. DOI: 10.3389/frobt.2019.00022.
- [15] M. Kasaei, K. K. Babarahmati, Z. Li, and M. Khadem, “A data-efficient neural ODE framework for optimal control of soft manipulators,” in *7th Annual Conference on Robot Learning*, 2023.
- [16] M. Hofer, L. Spannagl, and R. D’Andrea, “Iterative learning control for fast and accurate position tracking with an articulated soft robotic arm,” in *IROS 2019*, 2019. DOI: 10.1109/IROS40897.2019.8967636.
- [17] D. Schindele and H. Aschemann, “P-type ilc with phase lead compensation for a pneumatically driven parallel robot,” in *2012 American Control Conference (ACC)*, 2012. DOI: 10.1109/ACC.2012.6314641.
- [18] H. Ma, D. Büchler, B. Schölkopf, and M. Muehlebach, “A learning-based iterative control framework for controlling a robot arm with pneumatic artificial muscles,” in *Proceedings of Robotics: Science and Systems XVIII (R:SS 2022)*, 2022. DOI: 10.15607/RSS.2022.XVIII.029.
- [19] M. Deisenroth and C. Rasmussen, “PILCO: A Model-Based and Data-Efficient Approach to Policy Search,” in *ICML 2011*, 2011.
- [20] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *CoRR*, 2017. DOI: 10.48550/arXiv.1707.06347.
- [21] P. Bevanda, M. Beier, S. Heshmati-Alamdari, S. Sosnowski, and S. Hirche, “Towards Data-driven LQR with Koopmanizing Flows,” *IFAC-PapersOnLine*, 6th IFAC Conference, vol. 55, no. 15, 2022. DOI: 10.1016/j.ifacol.2022.07.601.
- [22] G. Torrente, E. Kaufmann, P. Foehn, and D. Scaramuzza, “Data-Driven MPC for Quadrotors,” *CoRR*, 2021. DOI: 10.48550/arXiv.2102.05773.
- [23] D. Bristow, M. Tharayil, and A. Alleyne, “A survey of iterative learning control,” *IEEE Control Systems Magazine*, vol. 26, no. 3, 2006. DOI: 10.1109/MCS.2006.1636313.
- [24] M. Meindl, D. Lehmann, and T. Seel, “Bridging reinforcement learning and iterative learning control,” *Frontiers in Robotics and AI*, vol. 9, 2022. DOI: 10.3389/frobt.2022.793512.
- [25] S. Bachhuber, I. Weygers, and T. Seel, “Neural ODEs for Data-Driven Automatic Self-Design of Finite-Time Output Feedback Control for Unknown Nonlinear Dynamics,” *IEEE Control Systems Letters*, vol. 7, 2023. DOI: 10.1109/LCSYS.2023.3293277.

APPENDIX

Algorithm 1 Input Trajectory Generator

Require: $T \in \mathbb{R}_{>0}$, $f \in \mathbb{N}_{>0}$

function GENERATESINUSOIDAL

$\omega \leftarrow 2\pi f$

$us \leftarrow \sin(\omega \text{range}(0, T, 0.01)) \frac{\sqrt{\omega}}{2}$

return us

Require: $T \in \mathbb{R}_{>0}$, $\Delta t_{\min} \in \mathbb{R}_{>0}$, $\Delta t_{\max} > \Delta t_{\min}$, $u_{\min/\max} \in \mathbb{R}$

function DRAWSPLINE

$ts, us \leftarrow [0.0], [0.0]$

while $\text{last}(ts) < T$ **do**

$t \leftarrow \text{last}(ts) + \text{uniform}(\Delta t_{\min}, \Delta t_{\max})$

$u \leftarrow \text{coinflip}(\text{last}(us), \text{uniform}(u_{\min}, u_{\max}))$

$ts, us \leftarrow \text{append}(ts, t), \text{append}(us, u)$

$us \leftarrow \text{cubicInterpolate}(\text{range}(0, T, 0.01), ts, us)$

return us