2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Design and Development of a Hydrogel-based Soft Sensor for Multi-Axis Force Control

Yichen Cai, David Hardman, Fumiya Iida, and Thomas George Thuruthel[1]

***Abstract*— As soft robotic systems become increasingly complex, there is a need to develop sensory systems which can provide rich state information to the robot for feedback control. Multi-axis force sensing and control is one of the less explored problems in this domain. There are numerous challenges in the development of a multi-axis soft sensor: from the design and fabrication to the data processing and modelling. This work presents the design and development of a novel multi-axis soft sensor using a gelatin-based ionic hydrogel and 3D printing technology. A learning-based modelling approach coupled with sensor redundancy is developed to model the environmentally dependent soft sensors. Numerous real-time experiments are conducted to test the performance of the sensor and its applicability in closed-loop control tasks at 20 Hz. Our results indicate that the soft sensor can predict force values and orientation angle within 4% and 7% of their total range, respectively.**

I. INTRODUCTION

Soft robotic sensing technologies are vital in providing dense tactile sensing abilities to emerging robotic systems in a conformable manner [1], [2]. These technologies allow us to measure the physical properties of interactions with minimal interference [3]. Numerous soft sensing technologies have been explored in the field recently, however, they have largely been limited to estimation of body configuration and contact localization [3], [2]. A largely unexplored sensory modality is multi-axis force sensing, a capability well-developed in rigid robotics [4]. This is primarily because of two reasons. First, for a soft bodied sensor, decoupling of multi-directional forces requires complex models and simplifying assumptions. This in turn makes the design and fabrication process challenging. Second, these stretchable materials exhibit several nonlinear material properties and are affected by environmental conditions, which present additional modelling challenges.

One of the first examples of a soft multi-axis force sensor used embedded microchannels filled with a conductive liquid which responded uniquely to different force directions applied to a force post [5]. This system could measure normal forces and shear force in a plane. However, there were still some challenges involved with the nonlinear properties of the sensor and its temperature dependence, which made it difficult to be used in real-world applications. Similarly, interlocking conductive microstructures have been used to develop sensors that can differentiate normal forces and shear, however, without any characterisation [6]. Similar to the work by Vogt et al.[5], magnetic field-based soft tri-axis tactile sensors were developed by Hongbo et al. using embedded hall effect sensors [7]. Along the same lines, capacitive [8] and inductive [2] elements can be used to differentiate multi-directional force components. More recently, magnetic skins have shown promise for multi-axis force sensing, especially over a larger area [9], [10]. Nevertheless, all these technologies still involve embedded rigid components, are highly location sensitive and have not yet been used for real-time control. Closing the loop requires accurate models of these sensors, which is challenging because of their nonlinear properties.

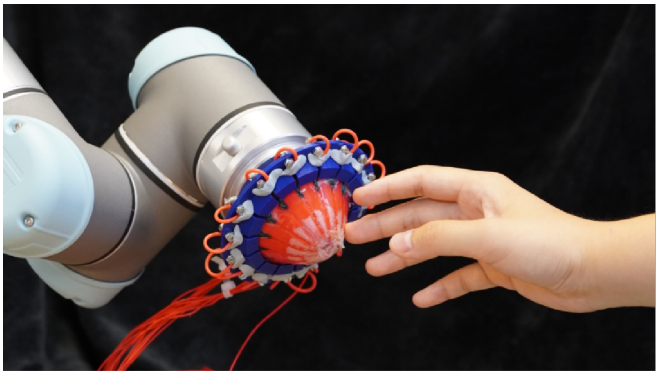

Fig. 1: The proposed three-axis soft sensor attached to the end-effector of a UR5 manipulator. The setup is trained to estimate the forces along the vertical axis, along with the direction of the applied force in two axes.

Gelatin-based hydrogels have recently been gaining interest in the field of soft robotic sensors because of their low cost, ease of fabrication, biodegradability, and printability [11], [12], [13]. Heiden et al. demonstrate their use as waveguides to facilitate basic directional sensing [14]. However, their highly linear resistive responses to applied strain [15] have yet to be implemented in a multi-axis sensor.

Time variant electro-mechanical properties are unavoidable in these sensor materials and their surrounding soft matrix. Hence, to model these soft strain sensors, the state estimation model must either use the time history of sensor states for future predictions [16] or use redundant sensor configurations to compensate for the nonlinearities [17]. Learning-based approaches using recurrent neural networks are typically used to directly incorporate the time dependent factors [18], [19], [20], [21]. Conversely, these time dependent hidden sensor states and external factors like temperature and humidity can be compensated using multiple

[1]All authors are with the Bio-Inspired Robotics Lab, University of Cambridge, UK. [yc539,dsh46,fi224,tg444]@cam.ac.uk

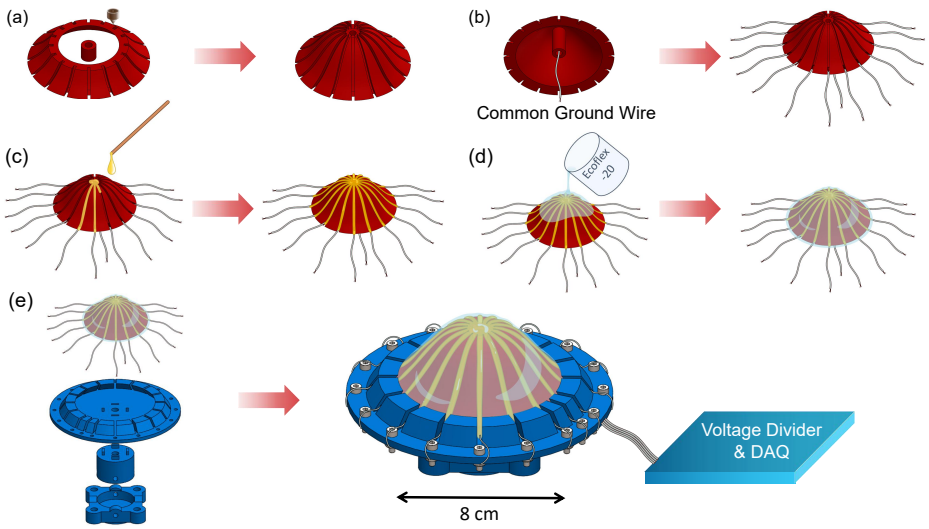

Fig. 2: Sensor fabrication process. (a) 3D-printing filaflex base (b) Attaching common ground wire and sensor wires (c) Applying salt-sensorized hydrogel to each groove (d) Spin-coat with Ecoflex-20 (e) Assemble to 3D-printed ABS base and wire to the voltage divider.

sensors [17], [15]. In this work, we use the combination of recurrent neural networks and redundant sensor configurations to model our soft sensors. Closing the loop with soft sensors can be done once a reliable state estimation model is created and the kinematic relation between estimated states and robot motion is established [22], [23], [24], [25].

This work presents the design, fabrication, modelling, and control of a fully soft multi-axis force sensor using an ionically conductive gelatin-glycerol hydrogel. Unlike existing works, our design is continuously soft and has a large sensing surface. The design is inexpensive and easy to fabricate, enabling quick prototyping and custom designs. To compensate for the hydrogel's baseline shift and environmental dependence, a modelling framework incorporating sensor redundancy and time-series learning is presented. Numerous experiments test the performance of the proposed sensor and its applicability for closed-loop control tasks.

II. Materials and Methods

A. Fabrication

The multi-axis force sensor is fabricated by 3D printing a cup-like, non-conducting flexible base (Filaflex) as mechanical support. The base is axially symmetric with 16 grooves extending from the centre to the edge, in which an ionically conductive gelatin-glycerol hydrogel is embedded (Figure 2). The hydrogel uses a 1:1.5:2.5:0.2:0.1 wt% composition of gelatin (pork, 240-260 bloom):glycerol:water:citric acid monohydrate: table salt (NaCl), as described in [15]. All grooves are connected and electrically grounded at the centre. The sensor is further coated with a layer of Ecoflex-20 to protect the hydrogel, then fastened to the end effector of a *Universal Robots* UR5 robotic arm using a base 3D printed from acrylonitrile butadiene styrene (ABS).

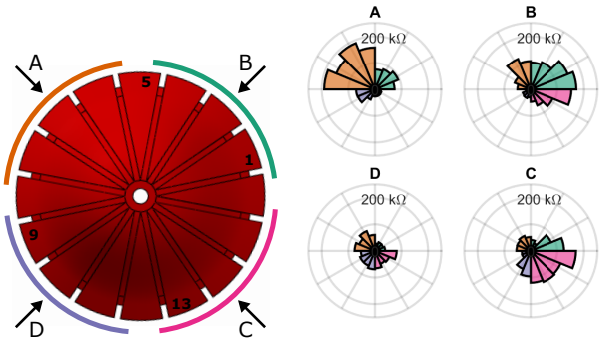

(a) (b)

Fig. 3: Resistive response of the 16 sensor channels to the 4 forces A-D. A video of this experiment can found in the supplementary material.

Each of the 16 hydrogel channels is connected through a voltage divider of resistance value 2.2 MΩ. It is then connected to a 16-bit National Instrument USB 6212 Series Data Acquisition System (DAQ), which supplies a 5V DC voltage across the divider and measures the output voltage at each central node at 20 Hz (Fig. 2). When an external force strains the hydrogel, its impedance and hence voltage values vary accordingly. The sensor morphology is designed such that the magnitude and direction of applied forces at the tip can be found by measuring the changes in the sensor properties. An illustration of the sensor response to forces from four different directions is shown in Figure 3. As designed, the sensor channels in each quadrant respond most strongly to the correspondingly directed forces, though with differing magnitude. The change in impedance is less

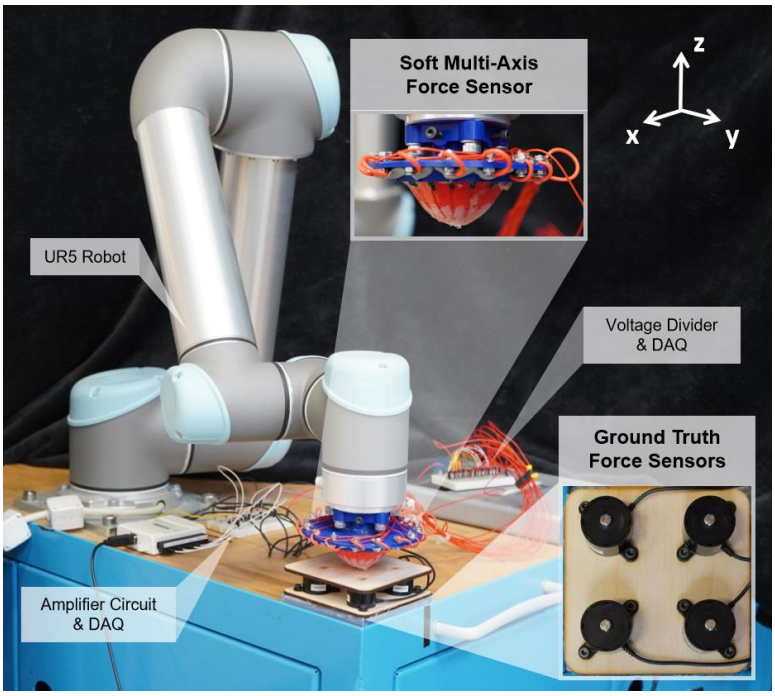

Fig. 4: Experiment setup. The ground truth force sensor consists of 4 load cells, where the signal is amplified by the amplifier circuit. The soft multi-axis force sensor is connected through a voltage divider.

obvious in direction D probably due to inconsistencies in the fabrication process. However, these asymmetric responses can easily be handled by any learning-based approach as long as the observed response is unique. Any environment dependent changes (e.g. temperature) to the hydrogel affects all the sensor channels similarly, and hence can be eliminated easily.

B. Experimental Setup

The soft multi-axis force sensor is installed at the tip of an UR5 robotic manipulator. For our experiments, the UR5 arm is constrained to rotational motions in the X and Y axes and translational motion in the Z-axis. To obtain ground truth force applied on the soft sensor, four FC22 compression load cells are placed in a square formation to support a rigid platform in the XY-plane (Figure 4). Raw signals from the FC22 load cells are amplified using Texas Instruments UA741CP operational amplifiers. These signals are read through analog channels of the DAQ, then read through USB by the MATLAB programming environment.

Estimating the magnitude and orientation of the exerted force from the sensor states can be a difficult problem due to multiple factors. First, because of the sensor's non-collocated architecture, the relationship between sensor response and the control variable is nontrivial, which in this case is the force in the Z-direction and the 2 DoF orientation information. Second, the sensor has significant temporal nonlinearities like drift and hysteresis. To address the first issue, we use multiple sensors and recurrent neural networks (RNNs) which utilize information from past data points for prediction [25].

C. Sensor Modeling

The relationship between the hydrogel sensor responses and the force output is modeled using a long-short term memory (LSTM) network. We collect 25 batches of data, each of varying time lengths of around 15 min, with 30-90s rest time between each batch to ensure that the initial condition varies. In each batch, the UR5 presses the sensor against the platform 31 times at random orientations, velocities, and for a random duration of time. Note that maximum rotation from the Z-axis is ±24.3 degrees for data collecting and all subsequent tests.

The orientation of the manipulator is updated at the start of each press before the sensor can make contact with the platform. The time between the orientation update and the first detection of exerted force is entirely random. Since the sensors respond to the orientation change once in contact, for training the ground truth orientation data is updated only when the force value goes above a small threshold. The aligned raw data is then separated into a training and testing set, each normalized by subtracting the set's initial value, then dividing by its standard deviation. This ensures that the sensors always start at zero to the network, even if the actual values drift over time.

A deep network consisting of two LSTM layers each with 20 hidden units and a dropout layer with dropout rate of 0.3 is used to map the 16 soft sensor inputs to the 4 force outputs and 3 orientation outputs (Figure 5). Of the 3 outputs, only the first 2 (representing rotations in the X and Y axes) are important since we are constraining the rotation in the z-axis, essentially controlling only two DoF orientations. Four batches out of the 25 from the start, middle, and end of the

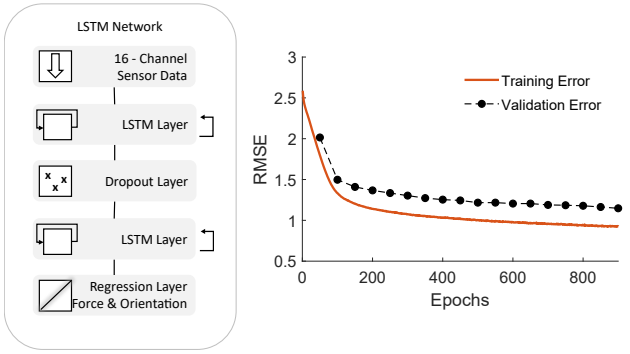

Fig. 5: The LSTM network used for estimating the force and orientation of contact from the raw sensor data (left). Training curve of the network (right).

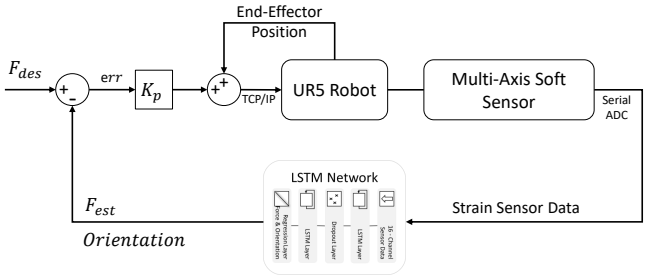

Fig. 6: Closed-loop control architecture using the learned state estimation model.

entire dataset are taken as validation to prevent overfitting. A mini-batch size of 500 is used for updating the weights of the network after each iteration. The MATLAB deep learning toolbox is used for creating and training the LSTM network. The adaptive moment estimation (Adam) optimizer is used to update the weights of the network. The root mean squared error (RMSE) for the training and validation data is shown in Figure 5.

D. Control

We use the LSTM network's predictions to design a proportional controller which runs at the same frequency (20Hz) as the learned model (Figure 6). A simple proportional feedback controller is used to close the loop. For our feedback experiments, we try to control the cumulative perpendicular force (F_{est}) applied by the manipulator irrespective of the orientation of the manipulator. During the test, the raw signal from the 16 sensors is measured, resampled to 20Hz by linear interpolation, normalized, and fed to the LSTM network in each control loop. The predicted forces values are added and fed to the control loop. The estimated orientation is not used for these control tasks. But note that indirectly, the network is using the orientation information to estimate the applied vertical force. The error between the estimated force and the target force is proportional to the distance that the robot is commanded to move, as shown in the equation below:

$$Translation = K_p(F_{des} - F_{est}) \quad (1)$$

This proportional gain is tuned manually in different control experiments to give the best results.

III. EXPERIMENTAL RESULTS

Three sets of experiments are performed to validate the proposed sensor and control framework. First, an open-loop test to validate the real-time accuracy of the learned model is performed. As this involves smooth motions with patterns similar to the training set, this result is a good indicator of the accuracy of the sensor system. Second, we analyze the closed-loop performance of the Z-axis force controller by performing a frequency sweep at a fixed random orientation. Finally, the steady-state performance of the closed-loop controller is evaluated for step inputs for randomly varying end-effector orientation.

The average and standard deviation values of the predicted-real and desired-real errors for each experiment are shown in Table I. The average predicted-real error is calculated by taking the absolute difference between predicted and real values, then taking the mean; similarly, the average desired-real error is the mean of the absolute difference between desired and real values, but only when the sensor is in contact with the ground truth force sensors. The real-time test does not specify desired force values, so it does not have an entry for desired-real force error. Steady-state and frequency sweep tests have fixed orientation, hence they do not have data for orientation error.

A. Real-time Open-loop Prediction

For this experiment, the sensor is rotated to a random orientation and pressed down on the ground truth force sensor at varying velocities and forces, and for a random duration of time. The maximum allowed rotation is ±24.3 deg from the z-axis. Figure 7 shows the real-time performance of the network on both the force and orientation prediction. The error of the orientation prediction is shown as the angle between the real and predicted Z vectors (Figure 7b). Error in the rotation about the z-axis is negligible since the sensor's actual and predicted rotation in the z-axis is extremely close to zero (of the order of 10^{-5}-10^{-3}). The predicted vertical forces have an average error of 4% of the total force range (Table I). There seems to be an onset of drift in the base value of the predictions over time indicating that the LSTM network has to be reset and the sensor states have to be reset to zero. The average error in orientation is slightly higher with an average error of 7% of the total range of rotation along the Z vector. Note that the orientation predictions are only output when the sensor is in contact with the environment. It can also be observed that there is a small remnant artifact even after contact removal either due to hysteresis or other factors. However, this is not found in the force predictions. Further studies have to be done to investigate the cause of this.

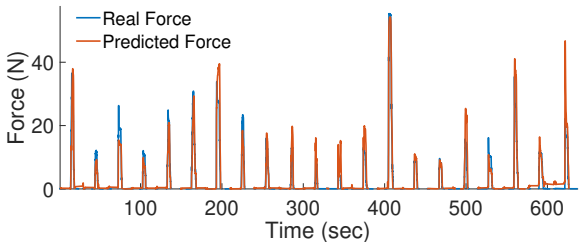

(a) Cumulative ground truth and predicted forces.

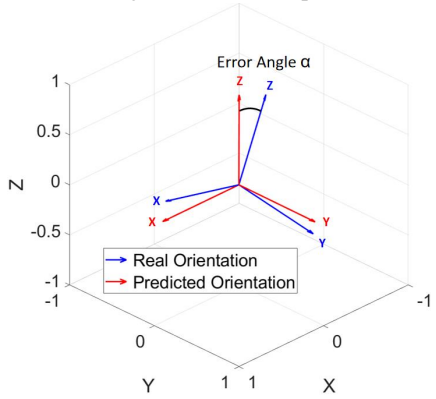

(b) Calculation of the orientation error.

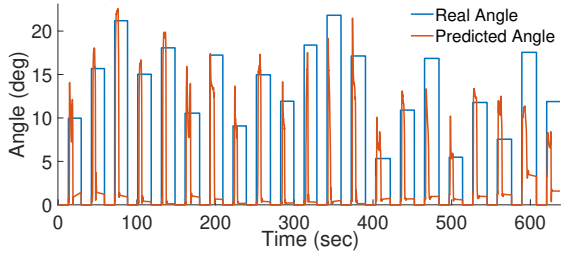

(c) Predicted rotation of the z vector and the actual rotation. Note that the orientation predictions are only output when the sensor is in contact with the environment.

Fig. 7: Real-time force and orientation prediction performance.

B. Closed-loop Frequency Sweep

For the next test, the closed-loop performance of the Z-axis force controller (See Figure 6) is tested by performing a frequency sweep. The target force is given as a chirp signal of increasing frequency, from 0Hz at t = 0s to 0.3Hz at t = 550s. The soft sensor is fixed at one orientation and only translational movement in the Z-axis is allowed. For this dynamic response test, the proportional gain factor (K_p) is tuned to be 0.4. This larger factor ensures that the robot moves quickly enough to follow the target force. The sensor's frequency response with target force as input and real force as output is shown in Figure 8. As the frequency increases linearly, however, the peak value of the ground truth force decreases, while the peak value of the predicted force drops by a small amount and then remains almost constant. This is expected because the proportional gain factor limits UR5's velocity and hence the maximum depth it can press down on the ground truth sensor. At higher frequencies the prediction error increases, indicating either conditions that have not been trained on or the limitations of the sampling rate. The performance of the sensor is robust up to the cutoff frequency at 0.12Hz (See Bode plot of the controller in Figure 9).

Predicted-Real Force Error		
Tests	Average Error (N)	Standard Deviation (N)
Real-time	1.9794	3.7260
Steady-state 0.1	5.5479	7.9809
Steady-state 0.4	6.0199	8.2739
Frequency Sweep	5.2102	5.5957

Desired-Real Force Error		
Tests	Average Error (N)	Standard Deviation (N)
Steady-state 0.1	16.4404	9.6384
Steady-state 0.4	9.6754	10.2800
Frequency Sweep	7.4863	9.4491

Orientation (Error Angle)		
Tests	Average Error (deg)	Standard Deviation (deg)
Real-time	1.6923	2.8449

TABLE I: Prediction performances for all three experiments.

C. Steady State Force Control

For the final test, the soft sensor is rotated to a random orientation, then pressed down onto the platform to exert a random predesignated constant target force for around 20s before it is removed. This procedure is repeated 21 times. The predicted force, real force, and target force values are plotted in Fig. 10(a). Two proportional gain factors K_p for the controller are tested: 0.1 and 0.4. The smaller gain factor successfully prevents the exerted force from overshooting when the UR5 presses down to meet the target force. However, it can be too slow to follow the force step input, leading to a large average error between real and desired force values. A larger gain factor of 0.4 makes sure that the predicted force quickly meets the target, though the real force may be much lower than predicted or may overshoot.

As the step input leads to an impulse impact, larger errors between the predicted and real forces are observed. This is in line with the observations from the frequency sweep, as the prediction errors start increasing at higher frequency. Adopting a faster sample rate and training the LSTM network on more dynamic motion data should improve performances. The latency between the real and target force is determined by the distance of the sensor from the surface and the gain factor. Shorter distance and larger gain factors would decrease the time to contact.

To show the possible applications of orientation prediction, we conduct two tests. Since the learned network is able to predict the sensor's orientation from its states when a force in the z-axis is applied, we are able to find the force vector given any sensor orientation when the force is exerted from a random direction. This is done by performing coordinate transformation based on UR5 kinematics. The force vector with respect to the UR5 base coordinate can be found and a closed-loop tactile follow-the-leader task can be performed. The supplementary video attachment shows this 3-axis feedback controller.

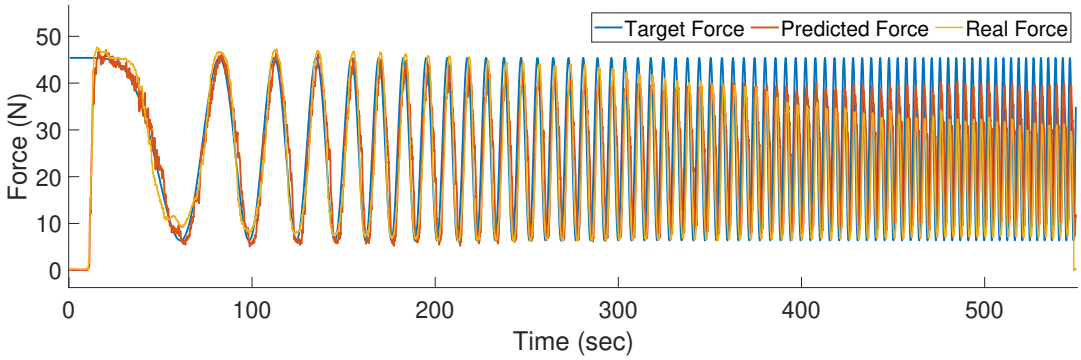

Fig. 8: Orientation invariant Z-axis frequency sweep.

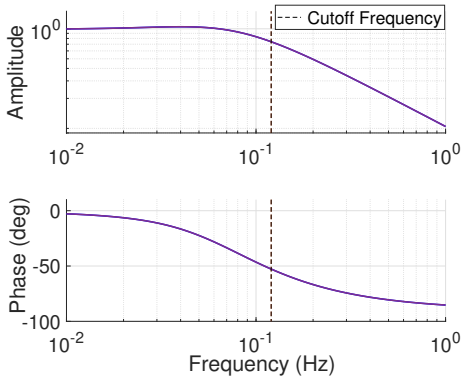

Fig. 9: Bode plot of the closed-loop controller.

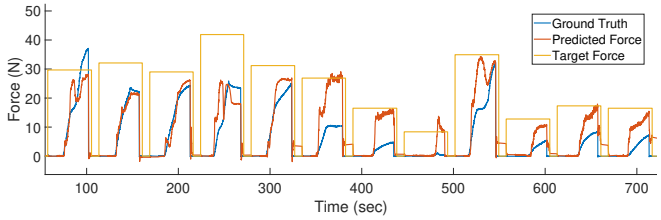

(a) K = 0.1

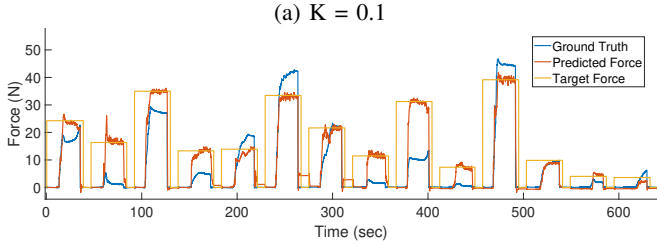

(b) K = 0.4

Fig. 10: Orientation invariant steady-state force control.

IV. CONCLUSION

This paper presents a fully soft gelatin-glycerol sensor for multi-axis force control. The radially symmetric design of the sensor allows us to differentiate forces and their magnitudes from multiple direction using minimal electrodes. The fabrication process uses 3D printing methods and casting. Using recurrent neural networks, the sensor is able to robustly predict the orientation and magnitude of the exerted force, even with significant drift and hysteresis in sensor response. This is because of the redundant sensor configuration and the dynamic modelling approach. With the trained network, a simple Proportional controller is designed for both static and dynamic tests. To the best of our knowledge, this is the first demonstration of a real-time multi-axis force controller using a fully soft sensing device.

Future works include better fabrication techniques for improving the robustness of the sensor and speed of production. Better encapsulation methods will be investigated to reduce direct contact of the hydrogel material to the environment. For example, the grooves could be switched to the inside surface of the Filaflex base so that the hydrogels will not be in direct contact with the environment, but will still deform with respect to applied forces. Different morphologies of the embedded hydrogels can also be investigated for varied functionalities. For example, the sensor grooves can be arranged into spiral shape for better torque estimation. Scaling up the system would involve using faster data acquisition systems with more sensory lines and using accurate multi-axis ground truth sensors. The current learning-based approach with recurrent neural networks to sensor modeling has been shown to be robust. However, network performance can be improved by collecting more data over a wide dynamic range, as well as more data across different environmental conditions. The feedback controller can be improved by adding a derivative component of the controller if noises in the predicted force can be reduced with additional data processing techniques.

ACKNOWLEDGMENT

This work was supported by the SHERO project, a Future and Emerging Technologies (FET) programme of the European Commission (grant agreement ID 828818), Mathworks Small Research Grant Programme, and by EPSRC DTP EP/R513180/1. For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising.

DATA AVAILABILITY

The code and data used for all reconstructions is hosted at https://github.com/Yichen-Cai-yc539/Design-and-Development-of-a-Hydrogel-based-Soft-Sensor-for-Multi-Axis-Force-Control.

REFERENCES

[1] B. Shih, D. Shah, J. Li, T. G. Thuruthel, Y.-L. Park, F. Iida, Z. Bao, R. Kramer-Bottiglio, and M. T. Tolley, "Electronic skins and machine learning for intelligent soft robots," *Science Robotics*, vol. 5, no. 41, 2020.

[2] H. Wang, M. Totaro, and L. Beccai, "Toward perceptive soft robots: Progress and challenges," *Advanced Science*, vol. 5, no. 9, p. 1800541, 2018.

[3] M. Amjadi, K.-U. Kyung, I. Park, and M. Sitti, "Stretchable, skin-mountable, and wearable strain sensors and their potential applications: a review," *Advanced Functional Materials*, vol. 26, no. 11, pp. 1678–1698, 2016.

[4] J. O. Templeman, B. B. Sheil, and T. Sun, "Multi-axis force sensors: A state-of-the-art review," *Sensors and Actuators A: Physical*, vol. 304, p. 111772, 2020.

[5] D. M. Vogt, Y.-L. Park, and R. J. Wood, "Design and characterization of a soft multi-axis force sensor using embedded microfluidic channels," *IEEE sensors Journal*, vol. 13, no. 10, pp. 4056–4064, 2013.

[6] J. Park, Y. Lee, J. Hong, Y. Lee, M. Ha, Y. Jung, H. Lim, S. Y. Kim, and H. Ko, "Tactile-direction-sensitive and stretchable electronic skins based on human-skin-inspired interlocked microstructures," *ACS nano*, vol. 8, no. 12, pp. 12 020–12 029, 2014.

[7] H. Wang, G. De Boer, J. Kow, A. Alazmani, M. Ghajari, R. Hewson, and P. Culmer, "Design methodology for magnetic field-based soft tri-axis tactile sensors," *Sensors*, vol. 16, no. 9, p. 1356, 2016.

[8] T. Nakadegawa, H. Ishizuka, and N. Miki, "Three-axis scanning force sensor with liquid metal electrodes," *Sensors and Actuators A: Physical*, vol. 264, pp. 260–267, 2017.

[9] Y. Yan, Z. Hu, Z. Yang, W. Yuan, C. Song, J. Pan, and Y. Shen, "Soft magnetic skin for super-resolution tactile sensing with force self-decoupling," *Science Robotics*, vol. 6, no. 51, p. eabc8801, 2021.

[10] T. Hellebrekers, O. Kroemer, and C. Majidi, "Soft magnetic skin for continuous deformation sensing," *Advanced Intelligent Systems*, vol. 1, no. 4, p. 1900025, 2019.

[11] M. Baumgartner, F. Hartmann, M. Drack, D. Preninger, D. Wirthl, R. Gerstmayr, L. Lehner, G. Mao, R. Pruckner, S. Demchyshyn, and Others, "Resilient yet entirely degradable gelatin-based biogels for soft robots and electronics," *Nature Materials*, pp. 1–8, 2020.

[12] J. Shintake, H. Sonar, E. Piskarev, J. Paik, and D. Floreano, "Soft pneumatic gelatin actuator for edible robotics," *arXiv*, pp. 6221–6226, 2017.

[13] B. Ying and X. Liu, "Skin-like hydrogel devices for wearable sensing, soft robotics and beyond," *Iscience*, vol. 24, no. 11, p. 103174, 2021.

[14] A. Heiden, D. Preninger, L. Lehner, M. Baumgartner, M. Drack, E. Woritzka, D. Schiller, R. Gerstmayr, F. Hartmann, and M. Kaltenbrunner, "3d printing of resilient biogels for omnidirectional and exteroceptive soft actuators," *Science Robotics*, vol. 7, no. 63, p. eabk2119, 2022.

[15] D. Hardman, T. George Thuruthel, and F. Iida, "Self-healing ionic gelatin/glycerol hydrogels for strain sensing applications," *NPG Asia Materials*, vol. 14, no. 1, pp. 1–13, 2022.

[16] T. G. Thuruthel, B. Shih, C. Laschi, and M. T. Tolley, "Soft robot perception using embedded soft sensors and recurrent neural networks," *Science Robotics*, vol. 4, no. 26, 2019.

[17] T. G. Thuruthel, J. Hughes, A. Georgopoulou, F. Clemens, and F. Iida, "Using redundant and disjoint time-variant soft robotic sensors for accurate static state estimation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2099–2105, 2021.

[18] D. Kim, J. Kwon, B. Jeon, and Y.-L. Park, "Adaptive calibration of soft sensors using optimal transportation transfer learning for mass production and long-term usage," *Advanced Intelligent Systems*, vol. 2, no. 6, p. 1900178, 2020.

[19] D. Kim, J. Kwon, S. Han, Y.-L. Park, and S. Jo, "Deep full-body motion network for a soft wearable motion sensing suit," *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 1, pp. 56–66, 2018.

[20] Z. Chen, R. Wu, S. Guo, X. Liu, H. Fu, X. Jin, and M. Liao, "3d upper body reconstruction with sparse soft sensors," *Soft Robotics*, 2020.

[21] R. L. Truby, C. Della Santina, and D. Rus, "Distributed proprioception of 3d configuration in soft, sensorized robots via deep learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3299–3306, 2020.

[22] J. C. Case, E. L. White, and R. K. Kramer, "Sensor enabled closed-loop bending control of soft beams," *Smart Materials and Structures*, vol. 25, no. 4, p. 045018, 2016.

[23] W. Felt, K. Y. Chin, and C. D. Remy, "Smart braid feedback for the closed-loop control of soft robotic systems," *Soft robotics*, vol. 4, no. 3, pp. 261–273, 2017.

[24] H. A. Sonar, A. P. Gerratt, S. P. Lacour, and J. Paik, "Closed-loop haptic feedback control using a self-sensing soft pneumatic actuator skin," *Soft robotics*, vol. 7, no. 1, pp. 22–29, 2020.

[25] T. George Thuruthel, P. Gardner, and F. Iida, "Closing the control loop with time-variant embedded soft sensors and recurrent neural networks," *Soft Robotics*, 2022.