Tactile perception in hydrogel-based robotic skins using data-driven electrical impedance tomography

David Hardman [a,*], Thomas George Thuruthel [a,b], Fumiya Iida [a]

[a] *Bio-Inspired Robotics Lab, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, UK*
[b] *Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, UK*

ARTICLE INFO

ABSTRACT

Combining functional soft materials with electrical impedance tomography is a promising method for developing continuum sensorized soft robotic skins with high resolutions. However, reconstructing the tactile stimuli from surface electrode measurements is a challenging ill-posed modelling problem, with FEM and analytic models facing a reality gap. To counter this, we propose and demonstrate a model-free superposition method which uses small amounts of real-world data to develop deformation maps of a soft robotic skin made from a self-healing ionically conductive hydrogel, the properties of which are affected by temperature, humidity, and damage. We demonstrate how this method outperforms a traditional neural network for small datasets, obtaining an average resolution of 12.1 mm over a 170 mm circular skin. Additionally, we explore how this resolution varies over a series of 15,000 consecutive presses, during which damages are continuously propagated. Finally, we demonstrate applications for functional robotic skins: damage detection/localization, environmental monitoring, and multi-touch recognition - all using the same sensing material.

1. Introduction

The sensorization of compliant functional materials is a vital step in the development of biomimetic skins for soft robots [1,2], requiring the development of multimodal sensing capabilities in flexible, stretchable, and healable materials. Recently, attention has shifted towards electrical impedance tomography (EIT) in soft sensor technologies for robotic applications [3,4], beyond its traditional use in the field of medical imaging [5]. EIT enables continuum sensing, few electrodes for high resolutions, and remote electrode positioning so as not to interfere with a robot's physical interactions. The combination of these advantages with advanced sensorized materials - such as self-healing hydrogels - would demonstrate a significant step forwards in the development of biomimetic skins, but their time-varying nonlinear material properties necessitate data-driven approaches [6,7]. In addition, traditional reconstruction techniques output a low-level conductivity maps, rather than the deformation state which is more useful for soft robotic skins [5,8].

Gelatin-glycerol hydrogels are of ever-increasing interest in soft robotics due to their printability, resilience, biodegradability, and excellent sensorization characteristics [7,9–13]. Whilst the fracture of ballistic gelatin has previously been explored using EIT methods [14], few works have combined these functional hydrogels with EIT-based soft sensing, due to the difficulties in tactile reconstruction. Zhang et al. [15] use a low conductivity gelatin-glycerol hydrogel sheet alongside

8-channel tomography to reconstruct highly conductive tactile stimuli, though this method does not extend to insulated touches. Ionically conductive approaches may serve to counter this: Zhao et al. [16] use an ionically conductive liquid and deformable elastic film to perform simple tactile reconstructions, with Soleimani and Friedrich [17] using a similar setup to detect forces as small as 64 mN. The combination of similar technologies with functional hydrogels can serve to advance the field of somatosensory robotic skins, provided the challenges in their reconstructions can be addressed.

When coupled with soft skins, EIT's governing equation may be expressed in the form of Poisson's equation [18,19]:

$$\nabla^2 v + \frac{1}{\sigma} \nabla \sigma \cdot \nabla v = 0 \tag{1}$$

where σ represents the conductivity field of a body (assumed isotropic and real), and v is the body's potential field. The forward problem - given a proposed conductivity field σ, find v - can be solved both numerically and, in certain cases, analytically [20–22]. Electrical impedance tomography concerns the inverse problem: finding σ from a selection of electrode v measurements and conservation-based boundary conditions. This problem is notoriously nonlinear and ill-posed [5,19,23], and considerable research has been devoted to its reconstruction techniques. Early attempts to simplify the problem use linearizations of Eq. (1) about a constant-conductivity condition [19,24]. Barber and Brown [18,23] analytically backproject the measured voltage gradients along

curved equipotentials. Further iterative algorithms apply Tikhonov & Landweber regularizations to the ill-posed problem [25–27]. Optimization approaches have become more viable with increased computational availability, aiming to propose σ fields which minimize the error between Eq. (1)'s proposed voltages and those measured at the electrodes. Genetic algorithms and particle swarm optimizations have been particularly studied for their ability to minimize this cost function [28–31].

The creation of EIT simulations & finite element models facilitates data-driven approaches to the conductivity predictions using neural networks [32–34]. Soft sensory implementations of the technology tend to be coupled with simulation-based training: Park et al. [3] perform real-time reconstruction of a hydrogel/silicone elastomeric skin using deep neural networks trained using finite element simulations, whilst Duan et al. use deep learning to remove unwanted artefacts in the deformations of a stretchable conductive fabric [35]. Data-driven approaches are not the only option for such sensors: Russo et al. [36] use Tikhonov regularization with a one-step Newton parameter optimization to reconstruct a conductive fabric's tactile stimuli. Though conductivity maps are sufficient for medical imaging, these soft sensory implementations must assume a close dependence between deformation and conductivity. Using a real-world data-driven approach, we instead aim to directly predict a soft skin's deformation state, without the need for simulations or a training period.

In this work, we estimate the deformation states of an ionically conductive gelatin-glycerol hydrogel membrane using small amounts of real-world data. Given the substrate-dependent effects which such a material introduces, a data-driven approach is employed: a robotic arm provides deformations of known location and depth, to which impedance responses around the skin's perimeter are measured and used as inputs. Given the large scale experimentation barrier to entry, this is - to the authors' knowledge - the first demonstration of any EIT sensorized skin which constructs models from physical data, without the need for simulations or analytic models. We illustrate its validity with demonstrations of high-resolution localizations and its robustness to small damages.

We first show how data-driven approaches can reduce the reality gap often faced by analytical methods or simulation-based training with functional materials [37]. As the models are developed from real-world data without implicit assumptions about the material properties, they outperform traditional analytical tools. We then propose our model-free weighted activation map (WAM) method, which requires fewer real-world data points for its reconstruction. Unlike analytical tools, the proposed method is generalizable to any sensing material, size and geometry, while compared to black-box learning-based approaches our methodology is highly data efficient and provides actual maps of the predicted deformations. Additionally, it is shown to localize contact significantly better than an equivalent neural network with small datasets of approximately 500 presses. We explore the ability of the proposed method in detecting and adapting to damages with samples of over 15,000 presses. Finally, we present the method's ability to monitor environmental changes, identify and localize damages, and respond to multiple contacts simultaneously; a capability not easily achievable with the existing learning frameworks.

2. Materials and methods

2.1. Sensorized hydrogel

The sensorized soft robotic skin is fabricated from a gelatin-glycerol hydrogel using 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), based on the work of Hardman et al. [7]. After being homogenized at a temperature of 50 °C, the mixture is manually cast into a laser-cut circular polymethyl methacrylate (PMMA) mold, with 180 mm inner diameter and 3 mm depth. The mold is left at room temperature for two days to ensure an equilibrium is reached with the environmental humidity. The evaporation of water during this process shrinks the skin's
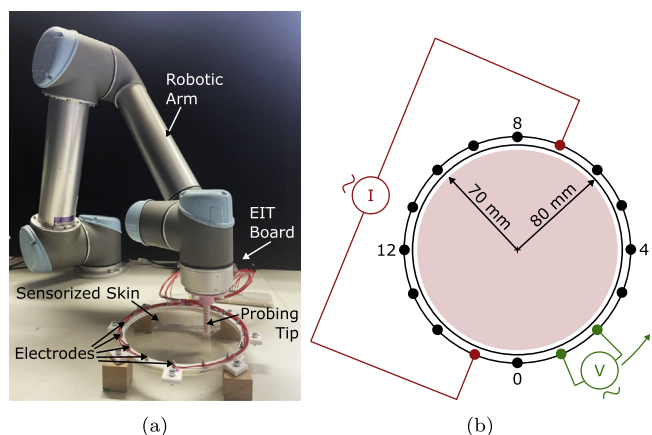
Fig. 1. a) Experimental setup: a sensorized hydrogel skin is probed by a robotic arm, whilst a series of measurements are performed by electrodes positioned around its perimeter. b) The skin's electrode configuration and shaded characterization area. Driving currents are applied to opposite electrodes, whilst the potential difference between all separate pairs of adjacent electrodes is measured.

diameter to approximately 160 mm, such that it is held under tension when secured in the 180 mm frame.

2.2. Hardware

Voltage measurements are collected using 16 electrodes arranged uniformly in a 170 mm diameter circle around a laser-cut PMMA frame with an inner diameter of 160 mm (Fig. 1). The frame consists of two halves, which are held tightly together with eight M4 fasteners. The sensorized skin (Section 2.1) is sandwiched between these two layers, pressing the 16 electrodes (M2 stainless steel bolts, with wires clamped between two nuts) firmly into its surface to ensure electrical contact. To prevent the skin from slipping within the frame, the two are joined between each electrode using a cyanoacrylate adhesive. The frame is supported 30 mm above a table by its perimeter (Fig. 1), so that any point on the exposed skin can be freely pressed: in this work, a maximum press depth of 20 mm is enforced.

The 16 electrodes are wired to a commercially available multi-channel *Spectra* electrical impedance measurement device (Mindseye Biomedical). For each measured state, the board uses a combination of 25 kHz opposite electrode constant driving currents to yield a set of 192 adjacent electrode voltage measurements [38]. The process is illustrated in Fig. 1: the driving current is sequentially applied to each of the $n = 16$ opposite electrode pairs, colored red. For each n, the potential differences between all $m = 14$ adjacent electrode pairs not sourcing or sinking current are measured (colored green). A USB connection returns the $n \times m$ measured impedances. By comparing the states before and after pressing, reconstruction algorithms (Section 2.3) aim to estimate the skin's conductivity field $\sigma(\mathbf{x})$.

Throughout the experiment, the skin is pressed using a Universal Robots UR5 robotic arm, equipped with a 3D printed polylactic acid (PLA) probe of 5 mm diameter and 40 mm length. The probe is covered with an insulating silicone bellow, and presses at random locations within a 140 mm diameter area of the skin, shaded red in Fig. 1. 15,000 such presses are recorded in succession, over the course of six days. One of four specified pressing depths - 5, 10, 15, or 20 mm - is randomly selected for each. Before each press, a full set of 192 measurements is recorded, and the probe descends with a maximum speed of 2 cm/s & acceleration of 2 cm/s^2. The press is held until a full set of 192 voltage measurements has been taken, before ascending with maximum speed and acceleration of 1 cm/s & 1 cm/s^2, respectively. A small strain-inflicted skin damage initiates between electrodes 2 & 3 close to
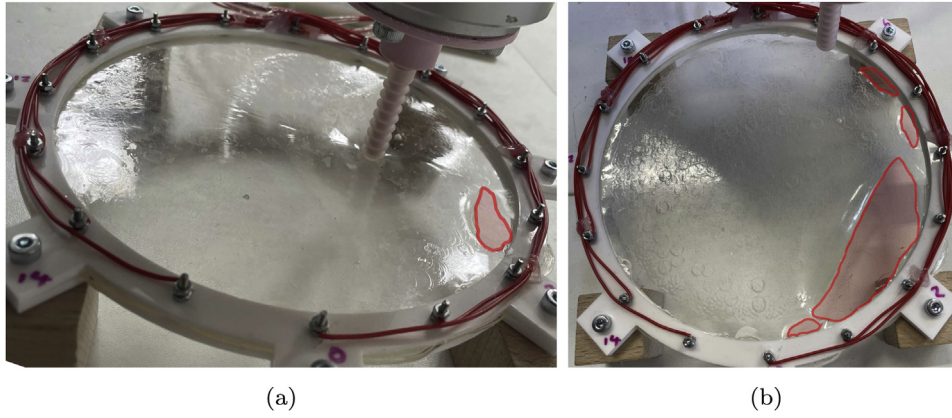
Fig. 2. Damage states after a) 5000 and b) 15,000 presses, where 10.7% of the surface is damaged. The damage initiates between electrodes 2 & 3 before propagating along the edge of the frame.

the frame's inner edge after approximately 5000 presses: Fig. 2a. As pressing continues, this damage propagates along the frame's edge, accelerating noticeably between 10,000 & 15,000 presses to reach 10.7% surface coverage. Figure 2b shows the damage in the skin's final state, with electrode 0 lowermost. It should be noted that no electrode is completely isolated by the damage; even in the final state, a small annular strip of the hydrogel runs within the enclosing frame.

A second undamaged skin is fabricated for the damage detection and multi-touch tests presented at the end of this work, constructed from 500 single press responses. Two probes are attached to the end-effector with a separation of 20, 75, or 108 mm, and are simultaneously pressed 10 mm into the skin's surface. Damages are inflicted using a craft knife, as shown in the supplementary video.

2.3. Reconstruction methods

Three reconstruction techniques are compared in this work: a model-based backprojection algorithm, which makes a number of assumptions about the underlying material properties; a simple data-driven feed-forward neural network, which can be trained to match the material's behaviours using thousands of observed responses; and the proposed superposition method using fewer observations - Weighted Activation Maps.

Weighted activation maps
The construction of Weighted Activation Maps (WAMs) is a data-driven approach which relies on the superposition of known responses and the incorporation of the known sensor geometry, significantly reducing the quantity of data required when compared to an equivalent neural network.

The maps are built from the responses to a set of 500 known presses. Each response consists of 192 values corresponding to the channel differences between the pressed and unpressed states. Each of the 192 channels is measured using a specific tetrapolar electrode configuration (Section 2.2). For each channel, an 'activation map' is created from the 500 contact locations: the 500 known responses ($\mathbf{v}$) are normalized and fed into a tanh function to eliminate outliers:

$$x_i = \tanh\left(\frac{v_i - \mu_\mathbf{v}}{\sigma_\mathbf{v}}\right). \tag{2}$$

Using the known positions of each press, a circular map of $\mathbf{x}$ is then linearly interpolated to visualize the regions in which the electrode configuration responds most strongly. The resolution of this visual map is dependent on the number of contact locations, and its accuracy is dependent on the uncertainty of the contact location (which is related to the accuracy of the robotic probe). Three such maps are shown in Fig. 3. It should be noted that the regions form paths which join their active

electrodes; each of the 192 paths thus has different patterns of sensitivity, which can be later superposed to predict the skin's deformation state.

Using the first 500 responses to random presses, 192 maps are built (Fig. 4). The randomized depths are not considered during the construction of the maps, to ensure noise-free localisation of both shallow and deep deformations. Given a response to an unknown press, the activation maps are used to predict its location using weighted superposition: each of the 192 maps is multiplied by its corresponding raw response value before all are added together to produce the final estimate. When predicting the value of a single press, the brightest pixel in the final image is used, illustrated by a black cross in Fig. 4.

Despite the simplicity of this approach, the rest of this work shows how the WAMs provide a valuable insight into the current paths through the conductive skin, achieving more accurate predictions than analytical models with less 'training' data than the alternate learning-based approach. The predictions remain robust as damage initiates and propagates, and the single-press data can be used to perform damage localization and multi-press predictions. From Fig. 4, we can see that the number of floating point operations required in a WAM reconstruction is equivalent to those encountered when passing through a 192×500 feedforward network layer. Its implementation is therefore expected to have a significantly quicker real-time refresh rate than a machine learning-based solution, which introduces significantly more layers and activation functions. In addition, once the data is available, the WAMs approach can immediately be deployed, whilst neural networks require an additional training period.

When visualizing the current paths, it should be noted that Fig. 3 applies an additional linear scaling with radius, since the responses near the edges tend to be greater than those at the centre. Equation (2) is replaced by:

$$p_i = \left| v_i\left(1.1 - \frac{r(v_i)}{R}\right) \right| \tag{3a}$$

$$x_i = \tanh\left(\frac{p_i - \mu_\mathbf{p}}{\sigma_\mathbf{p}}\right) \tag{3b}$$

This scaling is not used in the construction of WAMs, and only aids to visualize Fig. (3)'s activation maps.

Backprojection
As an example of a purely analytical reconstruction technique, we use a common standarized backprojection algorithm, implemented using *pyEIT* [39]. For each 192D response to a press, the algorithm aims to construct a 363-value conductivity map, visualized using *Matplotlib*'s tripcolor function. The brightest pixel in the resulting image is taken to be the method's prediction of press location.
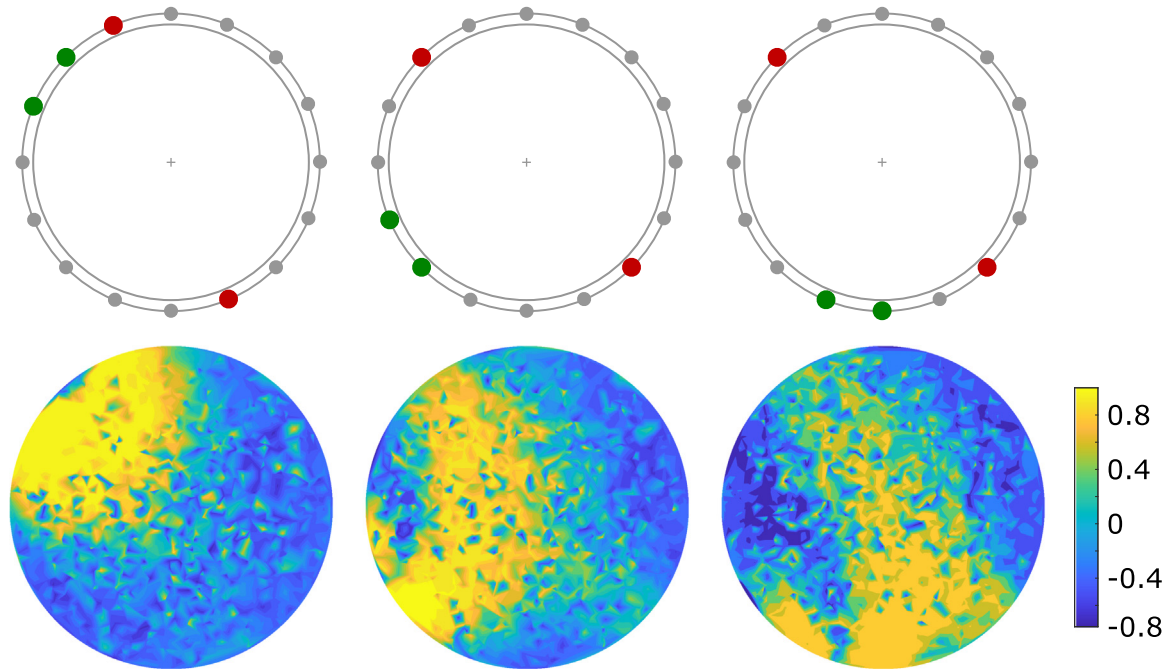
Fig. 3. Three exemplary activation maps, corresponding to channels 115, 130, & 132 during the first 500 random presses. The active tetrapolar electrodes are colored red (input) and green (output), between which current paths are clearly visible. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
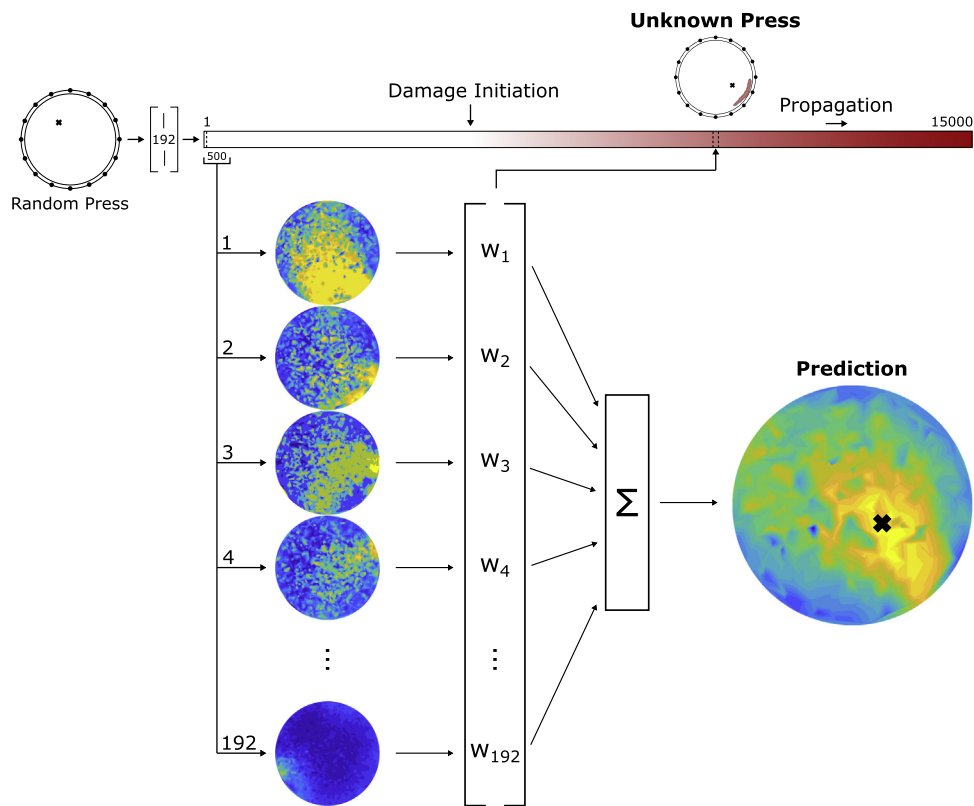

Fig. 4. Construction of a Weighted Activation Map (WAM). 192 activation maps are weighted using the response vector from an unknown press, and superposed into a final prediction image. The activation maps are created using data from the first 500 presses. The WAM method can continue to localize stimuli whilst damages are propagating.
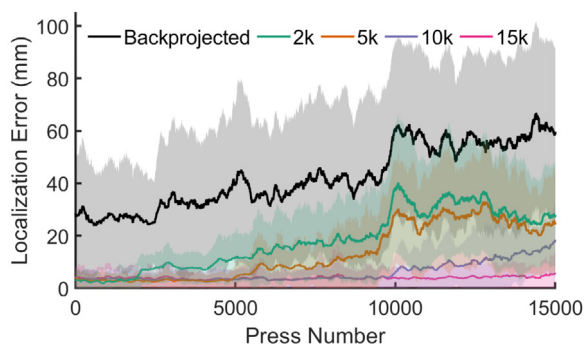
Fig. 5. A comparison of analytical backprojection and neural network (with different training sizes) localizations. Localization errors are plotted during the 15,000 presses during which damage initiates and propagates.

Table 1

Average resolutions (mm) and algorithmic complexities of the WAM, NN, and analytic approaches with small datasets. To match a typical application of sensor training and deployment, resolutions are calculated for the first 100 presses *after* the training window, rather than using an internal test set.

| Data Points (N)	100	500	1000
WAM Resolution (mm)	13.7	12.1	11.2
WAM Complexity (kflops)	38	192	383
NN Resolution (mm)	20.1	15.0	12.5
NN Complexity (kflops)	147*	147*	147*
Analytic Resolution (mm)	26.6	24.1	26.1
Analytic Complexity (kflops)	139	139	139

*Not including the network's training. Application of an activation function is assumed to be equivalent to 1 flop.

Neural networks

Feedforward neural networks are used to predict the location and depth of the presses, trained using only real-world data with a 80:10:10% training/validation/test data split. A fixed architecture is used throughout the work, implemented using *MATLAB*'s Deep Learning Toolbox. An input layer of size 192 is followed by three hidden layers with 300, 50, & 20 nodes respectively, and a 3-node regression output. A tanh activation layer follows each hidden layer. Training (stochastic gradient descent with momentum) begins with a learn rate of 0.05, which falls by 1% every 500 epochs, using minibatches of size 500. Training continues for 5000 epochs.

Before training, each input (taken to be the difference in response between the pressed and unpressed states in each of the 192 channels) is normalized to have zero mean and unity standard deviation, and each output is scaled to fall in the range (0, 1). These same normalization parameters are applied to all subsequent responses used for predictions.

3. Results

3.1. Localization

Before testing the WAM reconstructions, Fig. 5 demonstrates how data-driven methods are significantly advantageous over traditional analytical reconstruction techniques given the hydrogel's responses, using the two existing methods presented in Section 2.3. Backprojection uses only one response vector to make each prediction, whilst the neural networks are trained using an 80%:10%:10% split of the first 2000, 5000, 10,000, or 15,000 presses. Figure 5 plots the rolling localization error of the 5 cases over the 15,000 presses. It is immediately apparent that the networks outperform the analytical reconstructions by learning the patterns behind the non-modelled material properties, consistently producing lower errors. All plots trend upwards as damage initiates and propagates, with a sharp climb at ~10,000 presses corresponding to a sudden degradation. Unsurprisingly, each network performs best below its number of training presses, where 80% have been directly used in its training. Beyond this, errors steadily increase, suggesting a constant change in the response caused by the damage and hydrogel-environment equilibrium (addressed in Fig. 8a's discussion).

Given this inability of the purely analytical reconstructions to adapt to changing conditions, Fig. 6 focuses on the comparison of the two data-driven reconstruction techniques: neural networks and weighted activation maps. We hypothesise that WAM's knowledge of the sensor geometry allows the method to outperform neural networks with low amounts of data, which we test in Fig. 6a. Networks with the same architecture as Fig. 5 are trained using the first $N = 100, 500, 1000, 2000, 5000,$ or $10,000$ known presses and tested on all 15,000 presses, and WAM predictions use activation maps constructed from the same number of presses (rather than 500, as described in Section 2.3). As be-

fore, all methods perform best within their training region, then devolve as damage propagates, with a sharp increase in error around 10,000 presses. When $N = 100$ & $N = 500$, WAMs clearly outperform the networks, producing satisfactory solutions when there is insufficient data for the networks to fully converge. Table 1 shows a comparison of the two methods' average resolutions given small amounts of data: with just 100 inputs, the WAMs achieve an average resolution of 13.7 mm over the 170 mm skin, compared to NN's 20.1 mm. Both perform significantly better than the average analytic (backprojected) error for the same data window, which considers only the 192D response in its prediction. With very small datasets, the WAM location - predicted using the brightest pixel approach- is also affected by the separation of the training points, and resolutions could likely be improved further by instead identifying and localizing the centroid of a region of interest in the output map. The table also compares the number of floating point operations required for each prediction. Neural networks, backprojection, and WAM with $N = 500$ all have comparable complexities. However, the complexity of WAM varies linearly with N, resulting in a trade-off between resolution and complexity for which $N = 500$ produces a reasonable compromise.

As the quantity of available data increases, the two error sets in Fig. 6 become comparable ($N = 2000$) before the networks begin to outperform the WAMs due to the availability of damaged data. WAMs are therefore most advantageous when real-world data is expensive to collect, and accurate reconstructions are required from small amounts of data. $N = 500$ is selected for subsequent tests, since this is well within the range in which WAMs perform best. It also appears the WAMs could be relatively robust to damages which initiate after the training period - we later examine this effect in more detail.

Another significant advantage of WAMs over the localization networks is that they produce detailed deformation maps, providing more transparency to each prediction and enabling applications beyond localization (later discussed in Section 3.2). Figure 6b plots these deformation maps when $N = 500$ for the undamaged skin. As seen in Fig. 6a, both methods yield very small errors within the training region (0 presses). After 5000 presses the material properties have shifted, but accurate predictions can still be made - where WAM has lower average errors than the neural network (6 a). However, both methods appear significantly worse during the 5 mm presses, due to the lower response magnitude. In particular, the response of the perimeter-positioned electrodes is least when the shallow presses occur near the centre of the membrane, as in the first plot of the 5000 presses.

To test this, Fig. 7 compares the average errors with depth over test sets of 500 presses throughout the dataset. Indeed, in 6 of the 7 test sets, the lightest presses are associated with significantly higher localization errors than the others. The 10 mm presses follow, with the 15 & 20 mm presses having the lowest errors. The figure also examines how the error magnitudes are affected by damage: values increase with damage until the 10,000 press point, where they plateau. This effect can be seen in Fig. 7's lower plots: by the end, large errors stem from the same point/s
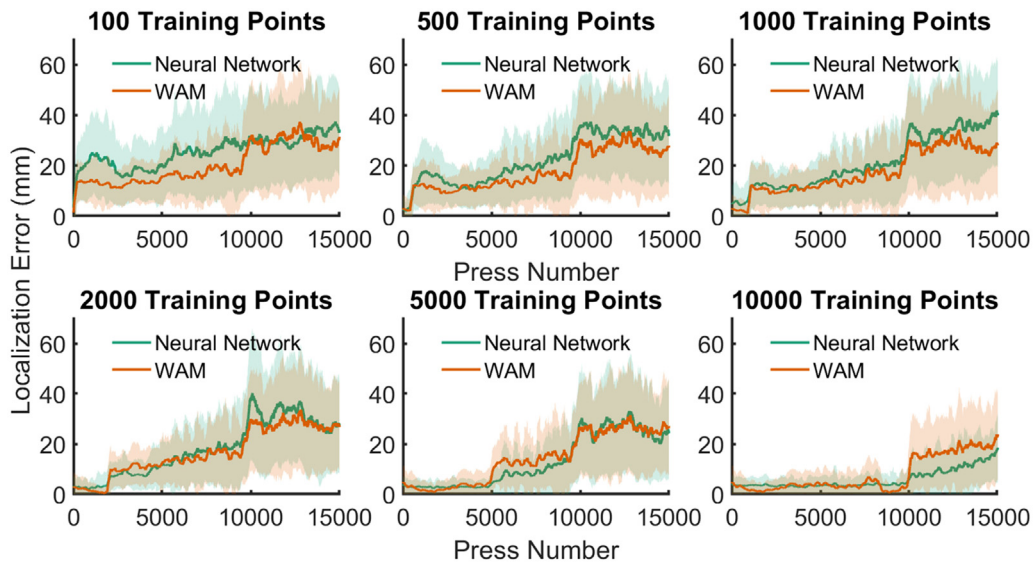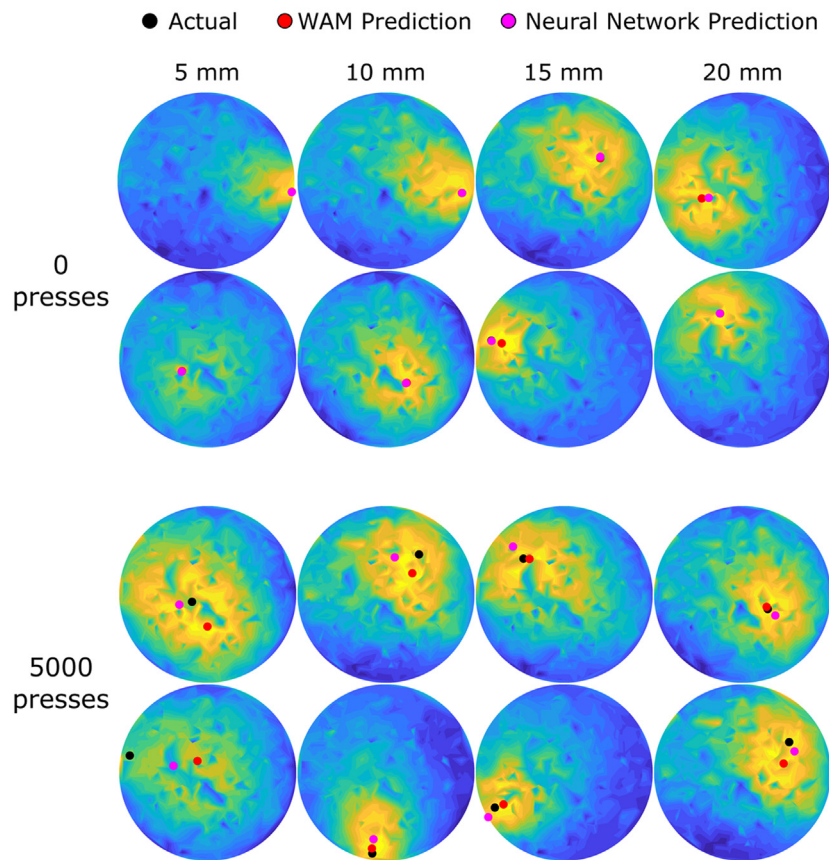
Fig. 6. a) Neural network and WAM localization errors when trained using different quantities of data. b) WAM visualizations and predictions for both methods when $N = 500$.

being repeatedly predicted. If the damage causes one channel of the response to rise significantly above the others, then this effect is expected. This could be countered by exploring additional normalization steps or tanh layers before the activation weightings are applied, minimizing the effect of outliers.

The only test set which does not follow these patterns is 1–500: the data on which the WAMs are constructed. All 4 errors are the lowest seen throughout the plot, but the errors *increase* with depth. No errors are noticeable in the 5 mm plots, since the small response magnitudes bias the predictions towards the individual activation maps created from
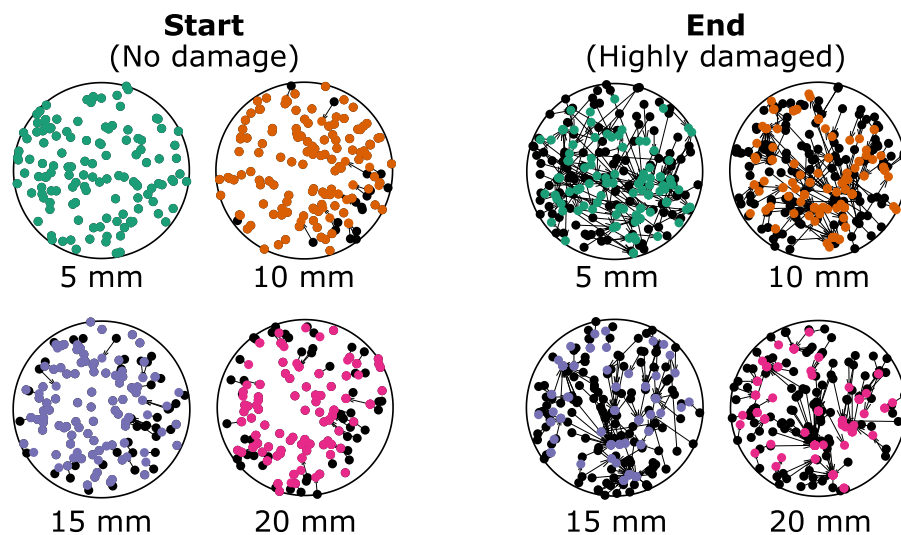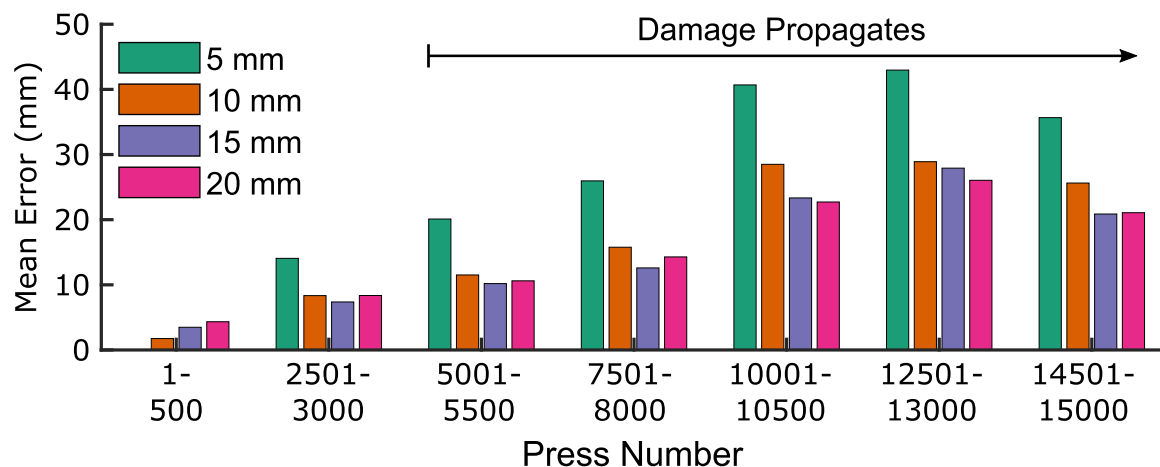
Fig. 7. Localization errors with time and depth of WAMs constructed with the first 500 presses. In the lower plots, black dots represent actual press locations, and colored dots represent WAM predictions.

the training data. The larger depths seem to favor predictions nearer the centre, underestimating the radial position of presses lying on the circle's edge. We hypothesize that this is due to deeper presses causing significant responses in multiple channels, adding more noise to the activation map selections and skewing the superposed predictions. However, further experimentation is required to verify and minimize this.

3.2. Extended applications

Having demonstrated WAMs' localization capabilities, Fig. 8 tests the 500 single press WAMs in a number of applications for which they have not been trained: environmental sensing, damage detection, and multi-touch localization. The central plot in Fig. 8a shows the response of the 192 channels to the unpressed state over the 6-day testing period. The hydrogel skin is known to respond to environmental fluctuations in temperature and humidity [7], which can be seen in the responses' 6 distinct regions. Before damage occurs (i.e. over the first ~5000 presses), these changes can be tracked using WAMs. By weighting the maps using the difference in unpressed response between a given press and the first press, the effects of these environmental fluctuations are visualized: at the peaks there is a clear response which sweeps across the centre of the hydrogel during homogenization (see accompanying video, which animates the fluctuations in baseline resistance during this testing period), whilst between these the skin periodically returns to a more homoge-

neous response. Though the capabilities of this behavior are not fully examined in this work, this preliminary demonstration illustrates the potential of the setup for inclusion in a multi-modal sensorized skin, where it can respond not only to tactile stimuli but also to moisture and temperature. When damage begins to propagate, these responses are dwarfed by the effect of the damage: the lower row of Fig. 8a plots the same difference in unpressed response for the later stages, where damage is significant. The damaged region responds more strongly over the propagation period, which suggests that single-press WAMs could be used for damage detection. This idea is explored in Fig. 8b, in which a new undamaged membrane is cast, mounted, and pressed 500 times to build the 192 required activation maps - this is done entirely autonomously using a robotic arm, enabling the quick and minimal-effort calibration of new skin materials and geometries. Rather than relying on reconstructed magnitudes to identify the occurrence of damages, continuously monitoring the unpressed baseline enables damages to be more easily identified: in the plot of the raw channel values, 3 scalpel cuts correspond to definite peaks in the measured response. By constructing a WAM from the difference in response before and after these peaks, the damage can be localized as well as identified. For the purposes of comparison, plot A visualizes the WAM when no damage has occurred - the corresponding time window is marked on the plot. Plot B shows the difference before/after the first small cut, made by pressing a scalpel into the surface. This allows the two halves to regain contact when the scalpel is removed, giving very little lasting change in the responses.
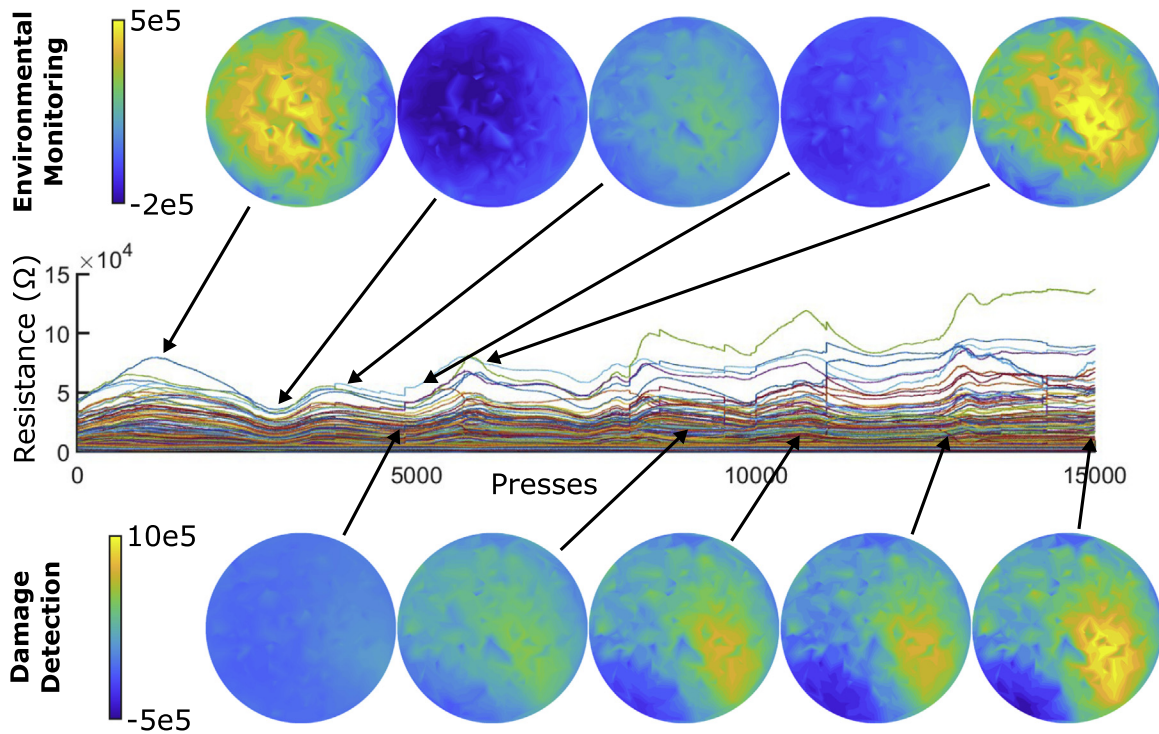
(a) Shifting baseline visualizations during environmental fluctuations and propagating damage.
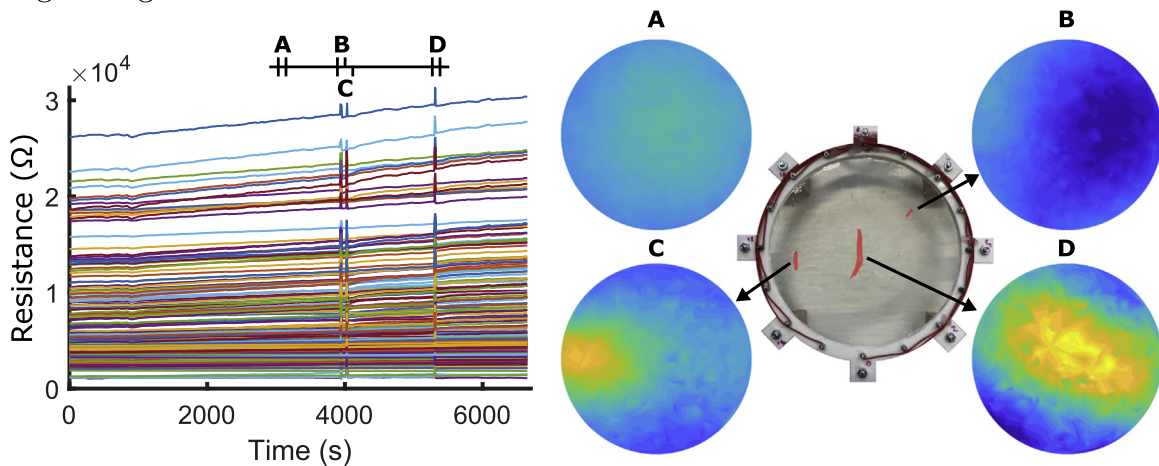
(b) Damage detection with a new skin.
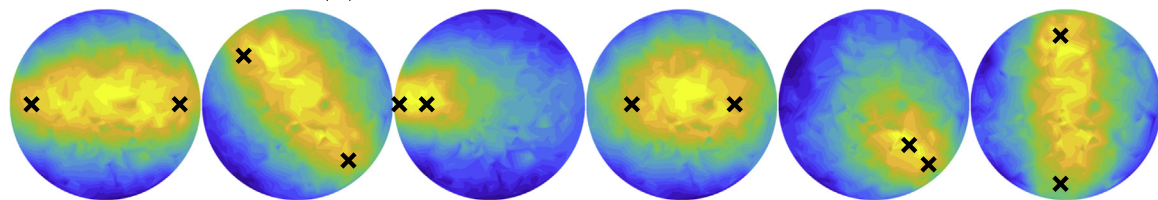
(c) Multi-touch localization.

Fig. 8. WAM applications beyond single-press localizations, despite this being the only data it has ever seen. a) Environmental monitoring and damage detection over 15,000 presses, constructed from the first 500 presses. The skin's state fluctuates with temperature and humidity. b) Damage identification and localization in a new membrane, using 500 single-presses. c) Multi-touch: WAM constructions for two simultaneous 10 mm presses at a range of separations.

Given the hydrogel's healability, these halves would be expected to heal over time if they remained in contact.

Conversely, damage C is large enough to yield a gap in the membrane, the localization of which is clearly visible. After allowing time for any transient responses to settle, the significant damage D is performed vertically in the membrane's centre. This causes a large response in the centre of the WAM. Though the response is in the correct region, it is not aligned perfectly with the cut (as seen in Fig. 8a). This is likely due to a breakdown of the linearity assumptions on which the WAMs are based, since the large damaged zone now returns a zero response.

Finally, Fig. 8c tests the WAM response to multiple touches, based on the single press constructions. Even at the 10 mm pressing depth, the maps show excellent matches to the probed locations, which are marked with black crosses. The maps highlight the entire region between the two electrodes, rather than two clear areas. As can be seen in the supplementary video, this reflects the membrane's tension increase along the entire line. As such, the weighted activation maps seem to be a good reflection of Eq. (1)'s reconstruction problem. The WAMs are maps of deformation rather than conductivity, aiding straightforward interpretation of the tactile signals. This technique helps to further data-driven proprioceptive & exteroceptive sensing in soft robotics, with our remote electrode configuration enabling continuous sensing of the skin without disturbing its properties and interactions.

4. Conclusion

In this work, we have demonstrated the data-driven tactile reconstructions of a gelatin-glycerol hydrogel skin, without the need for simulations or analytic modelling. We proposed the WAM - weighted activation maps - method, which is able to outperform learning-based methods for small datasets up to 1000 responses, whilst also providing a visualization of the predicted deformation state. By focusing on the skin's deformation state, the method is able to be directly applied to new tasks without retraining, including environmental monitoring, damage detection and localization, and multi-touch recognition. This discovery is particularly useful in the implementation of soft EIT skins in robotic systems, where analytic assumptions may be invalid, simulations may face the reality gap, and data may be expensive to collect. The perimeter-positioned electrodes suit the development of soft robotic skins using this technology, and future work will consider the sensorization and self-supervised learning of an entire hand, using electrodes placed only at its wrist. Finally, the preliminary results presented in Fig. 8 demonstrate the potential of further investigations into WAMs' use in shape recognition, multi-touch responses, and environmental monitoring. The challenges and methods associated with each of these will be developed in future works: finding the damage extent at which WAMs' assumptions become invalid; extracting multi-touch predictions from the reconstructed deformation maps; and combining damage monitoring with active healing techniques to create a closed-loop protection system. Additionally, the transferabilities of the activation maps between materials and geometries will be investigated. By exploiting similarities, the amount of data required for calibration could be reduced to even lower levels, facilitating quick deployment of custom sensorized skins.

Data Availability

The code and data used for all reconstructions is hosted at: https://github.com/DSHardman/WeightedActivationMaps.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.mtelec.2023.100032.

References

[1] B. Shih, D. Shah, J. Li, T.G. Thuruthel, Y.-L. Park, F. Iida, Z. Bao, R. Kramer-Bottiglio, M.T. Tolley, Electronic skins and machine learning for intelligent soft robots, Sci. Robot 5 (2020) 9239.

[2] F. Liu, S. Deswal, A. Christou, Y. Sandamirskaya, M. Kaboli, R. Dahiya, Neuro-inspired electronic skin for robots, Sci. Robot 7 (2022) 7344.

[3] K. Park, H. Yuk, M. Yang, J. Cho, H. Lee, J. Kim, A biomimetic elastomeric robot skin using electrical impedance and acoustic tomography for tactile sensing, Sci. Robot 7 (2022) 7187.

[4] K. Liu, Y. Wu, S. Wang, H. Wang, H. Chen, B. Chen, J. Yao, Artificial sensitive skin for robotics based on electrical impedance tomography, Adv. Intell. Syst. 2 (2020) 1900161, doi:10.1002/aisy.201900161.

[5] T.A. Khan, S.H. Ling, Review on electrical impedance tomography: artificial intelligence methods and its applications, Algorithms 12 (2019), doi:10.3390/a12050088.

[6] S. Terryn, D. Hardman, T.G. Thuruthel, E. Roels, F. Sahraeeazartamar, F. Iida, Learning-based damage recovery for healable soft electronic skins, Adv. Intell. Syst. (2022) 2200115, doi:10.1002/aisy.202200115.

[7] D. Hardman, T. George Thuruthel, F. Iida, Self-healing ionic gelatin/glycerol hydrogels for strain sensing applications, NPG Asia Mater. 14 (1) (2022).

[8] W. Xin, F. Zhu, P. Wang, Z. Xie, Z. Tang, C. Laschi, Electrical impedance tomographic shape sensing for soft robots, IEEE Rob. Autom. Lett. (2023).

[9] A. Heiden, D. Preninger, L. Lehner, M. Baumgartner, M. Drack, E. Woritzka, D. Schiller, R. Gerstmayr, F. Hartmann, M. Kaltenbrunner, 3D printing of resilient biogels for omnidirectional and exteroceptive soft actuators, Sci. Rob. 7 (63) (2022) eabk2119.

[10] M. Baumgartner, F. Hartmann, M. Drack, D. Preninger, D. Wirthl, R. Gerstmayr, L. Lehner, G. Mao, R. Pruckner, S. Demchyshyn, Others, Resilient yet entirely degradable gelatin-based biogels for soft robots and electronics, Nat. Mater. (2020) 1–8.

[11] J. Shintake, H. Sonar, E. Piskarev, J. Paik, D. Floreano, Soft pneumatic gelatin actuator for edible robotics, arXiv (2017) 6221–6226.

[12] R. Dai, H. Zhou, W. Huang, C. Li, C. Qin, X. Liu, Z. Pan, Conductive hydrogel-based electronics for intelligent sensing and smart controlling, J. Nanoelectron. Optoelectron. 16 (5) (2021) 689–698.

[13] W. Lee, V. Lee, S. Polio, P. Keegan, J.-H. Lee, K. Fischer, J.-K. Park, S.-S. Yoo, On-demand three-dimensional freeform fabrication of multi-layered hydrogel scaffold with fluidic channels, Biotechnol. Bioeng. 105 (6) (2010) 1178–1186.

[14] E.G. Lucas, Use of an electrical impedance tomography method to detect and track fractures in a gelatin medium, Master's Thesis, 2014, 10.37099/mtu.dc.etds/742

[15] H. Zhang, A. Kalra, A. Lowe, Y. Yu, G. Anand, A hydrogel-based electronic skin for touch detection using electrical impedance tomography, Sensors 23 (3) (2023), doi:10.3390/s23031571.

[16] T. Zhao, C. Wu, M. Soleimani, Ionic liquid based distributed touch sensor using electrical impedance tomography, IOP SciNotes 1 (2) (2020) 025005, doi:10.1088/2633-1357/abb345.

[17] M. Soleimani, M. Friedrich, E-skin using fringing field electrical impedance tomography with an ionic liquid domain, Sensors 22 (13) (2022), doi:10.3390/s22135040.

[18] D.C. Barber, B.H. Brown, Applied potential tomography, J. Phys. E Sci. Instrum. 17 (1984).

[19] M. Cheney, D. Isaacson, J.C. Newell, Electrical impedance tomography*, Soc. Ind. Appl. Math. 41 (1999) 85–101.

[20] S. Rush, Methods of measuring the resistivities of anisotropic conducting media in situ, J. Res. Natl. Bur. Stand. (1934) 66 (1962).

[21] P.W. Murray, Field calculations in the head of a newborn infant and their application to the interpretation of transcephalic impedance measurements, Biol. Eng. Comput 19 (1981) 546.

[22] K.A. Dines, R.J. Lytle, Analysis of electrical conductivity imaging, Geophysics 46 (1981) 1025–1036.

[23] F. Santosa, M. Vogelius, A backprojection algorithm for electrical impedance imaging, SIAM J. Appl. Math. 50 (1990) 216–243.

[24] Y. Yamashita, T. Takahashi, Methods and feasibility of estimating impedance distribution in the human torso, in: N. Svartholm (Ed.), Proc. VIth Int. Conf. Electrical Bioimpedance (ICEBI), 1981, pp. 87–90.

[25] Tikhonov regularization and prior information in electrical impedance tomography, IEEE Trans. Med. Imaging 17 (1998) 285–293, doi:10.1109/42.700740.

[26] Isaksen, A review of reconstruction techniques for capacitance tomography, Meas. Sci. Technol. 7 (1996) 325–337.

[27] H. Wang, C. Wang, W. Yin, A pre-iteration method for the inverse problem in electrical impedance tomography, IEEE Trans. Instrum. Meas. 53 (2004) 1093–1096, doi:10.1109/TIM.2004.831180.

[28] R. Olmi, M. Bini, S. Priori, A genetic algorithm approach to image reconstruction in electrical impedance tomography, IEEE Trans. Evol. Comput. 4 (2000) 83–88, doi:10.1109/4235.843497.

[29] G. Zheng, D. Chen, X. Yu, in: Image reconstruction algorithm for two-phase flow electrical capacitance tomography system, Vol. 4902, SPIE, 2002, p. 588, doi:10.1117/12.467340.

[30] P. Wang, J.S. Lin, M. Wang, An image reconstruction algorithm for electrical capacitance tomography based on simulated annealing particle swarm optimization, J. Appl. Res. Technol. 13 (2015a) 197–204.

[31] H. Wang, G. Xu, S. Zhang, W. Yan, Optimized excitation mode for generalized back projection algorithm in 3-D EIT, IEEE Trans. Magn. 51 (2015b), doi:10.1109/TMAG.2014.2361771.

[32] J. Li, X. Yang, Y. Wang, R. Pan, in: An image reconstruction algorithm based on RBF neural network for electrical capacitance tomography, 2012, doi:10.1109/ICEF.2012.6310416.

[33] A benchmark dataset and deep learning-based image reconstruction for electrical capacitance tomography, Sensors (Switzerland) 18 (2018), doi:10.3390/s18113701.

[34] M. Stasiak, J. Sikora, S.F. Filipowicz, K. Nita, Principal component analysis and artificial neural network approach to electrical impedance tomography problems ap-proximated by multi-region boundary element method, Eng. Anal. Bound. Elem. 31 (2007) 713–720, doi:10.1016/j.enganabound.2006.12.003.

[35] X. Duan, S. Taurand, M. Soleimani, Artificial skin through super-sensing method and electrical impedance data from conductive fabric with aid of deep learning, Sci. Rep. 9 (2019), doi:10.1038/s41598-019-45484-6.

[36] S. Russo, S. Nefti-Meziani, N. Carbonaro, A. Tognetti, A quantitative evaluation of drive pattern selection for optimizing EIT-based stretchable sensors, Sensors (Switzerland) 17 (2017), doi:10.3390/s17091999.

[37] G. Wang, J.C. Ye, B. De Man, Deep learning for tomographic image reconstruction, Nat. Mach. Intell. 2 (12) (2020) 737–748.

[38] A. Adler, P.O. Gaggero, Y. Maimaitijiang, in: Adjacent Stimulation and Measurement Patterns Considered Harmful, Vol. 32, IOP Publishing Ltd, 2011, pp. 731–744, doi:10.1088/0967-3334/32/7/S01.

[39] B. Liu, B. Yang, C. Xu, J. Xia, M. Dai, Z. Ji, F. You, X. Dong, X. Shi, F. Fu, pyEIT: a python based framework for electrical impedance tomography, SoftwareX 7 (2018) 304–308, doi:10.1016/j.softx.2018.09.005.