

# Predicting the Force Map of an ERT-Based Tactile Sensor Using Simulation and Deep Networks

Hyosang Lee<sup>1</sup>, Huanbo Sun<sup>1</sup>, Hyunkyu Park<sup>1</sup>, Gokhan Serhat<sup>2</sup>, *Member, IEEE*,  
Bernard Javot<sup>3</sup>, *Member, IEEE*, Georg Martius<sup>4</sup>, and Katherine J. Kuchenbecker<sup>5</sup>, *Fellow, IEEE*

**Abstract**—Electrical resistance tomography (ERT) can be used to create large-scale soft tactile sensors that are flexible and robust. Good performance requires a fast and accurate mapping from the sensor's sequential voltage measurements to the distribution of force across its surface. However, particularly with multiple contacts, this task is challenging for both previously developed approaches: physics-based modeling and end-to-end data-driven learning. Some promising results were recently achieved using sim-to-real transfer learning, but estimating multiple contact locations and accurate contact forces remains difficult because simulations tend to be less accurate with a high number of contact locations and/or high force. This paper introduces a modular hybrid method that combines simulation data synthesized from an electromechanical finite element model with real measurements collected from a new ERT-based tactile sensor. We use about 290 000 simulated and 90 000 real measurements to train two deep neural networks: the first (Transfer-Net) captures the inevitable gap between simulation and reality, and the second (Recon-Net) reconstructs contact forces from voltage measurements. The number of contacts, contact locations, force magnitudes, and contact diameters are evaluated for a manually collected multi-contact dataset of 150 measurements. Our modular pipeline's results outperform predictions by both a physics-based model and end-to-end learning.

**Note to Practitioners**—ERT-based tactile sensors use high-speed voltage measurements from electrodes distributed over a piezoresistive area to output a force map that shows where contact is occurring, and how strong each contact is. Such sensors

hold promise for giving robots and other devices a sense of touch over large surfaces with low hardware complexity. However, the software problem of converting voltages to an accurate force map has not previously been solved well, requiring either extensive model calibration or extensive contact data collection. This paper suggests a hybrid approach where a straightforward physics model simulates multi-contact experiments that are too costly to acquire in reality and a practical automatic indentation setup acquires real but geometrically limited multi-contact data. Although the number of real measurements required to learn the discrepancy between the sensor and the model is still large due to the inherent inverse nature of ERT-based tactile sensors, our combination of simulation and deep networks achieves better performance than either physical modeling or learning alone. This approach can advance practical large-area tactile sensing for industrial automation systems where multiple contacts occur, such as in manufacturing and assistive robotics. It could also likely be adapted to other nonlinear inverse problems.

**Index Terms**—Multi-point tactile sensing, electrical impedance tomography, electromechanical modeling, sim-to-real transfer, machine learning.

## I. INTRODUCTION

**T**ACTILE sensing is an important modality for autonomous robots to experience physical interactions. The development of suitable sensors for large surfaces has lately gained increasing research attention. Robotic tactile skin should not only be able to sense multiple simultaneous contacts but also scale to cover sensing areas that are large and often curved. This requirement has prompted most large-scale tactile sensing studies to utilize many small sensing elements in the form of arrays [1] or modules connected by digital communication [2]–[4]. These approaches can provide high spatial resolution and sensitivity if the sensing elements are located in a compact region. However, implementing such a system on a large and curved three-dimensional surface is challenging in terms of cost-efficiency, robustness, and data communication complexity [5].

Reconstruction approaches have emerged as a promising alternative means to create large-area tactile sensors. These methods use computation to estimate physical contact information over the sensor from a small number of indirect measurements. The measurements can be taken by various sensors such as optical cameras with markers [6], photodiodes and emitters [7], [8], pressure sensors covered with an elastic layer [9], [10], sparsely located strain gauges [11], [12], and resistance measurements [13]–[18]. Among these methods, resistance measurement is particularly favorable for

Manuscript received 14 September 2021; revised 17 January 2022; accepted 27 February 2022. Date of publication 17 March 2022; date of current version 6 January 2023. This article was recommended for publication by Associate Editor Q. Xu and Editor L. Zhang upon evaluation of the reviewers' comments. This work was supported in part by the Max Planck Society in Germany. The work of Huanbo Sun was supported in part by the China Scholarship Council and in part by the Max Planck Society. The work of Georg Martius was supported in part by the Machine Learning Cluster of Excellence, EXC Number 2064/1, under Project 390727645; and in part by the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ) under Grant 01IS18039B. The work of Katherine J. Kuchenbecker was supported by the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ) under Grant 01IS18039B. (*Corresponding author: Hyosang Lee.*)

Hyosang Lee, Gokhan Serhat, Bernard Javot, and Katherine J. Kuchenbecker are with the Haptic Intelligence Department, Max Planck Institute for Intelligent Systems, 70569 Stuttgart, Germany (e-mail: hslee@is.mpg.de).

Huanbo Sun and Georg Martius are with the Autonomous Learning Group, Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany.

Hyunkyu Park is with the Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology, Daejeon 34141, South Korea.

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/TASE.2022.3156184>.

Digital Object Identifier 10.1109/TASE.2022.3156184

whole-body tactile sensing because piezoresistive skins can be soft, thin, and fabricated in a wide variety of shapes.

As a well-established resistance reconstruction technique, electrical resistance tomography (ERT) [19] is widely adopted in large-area tactile sensing to estimate physical contact information using a small number of spatially distributed electrodes [13]–[18]. The electrodes are used to inject current into the piezoresistive sensing surface and measure the corresponding voltages formed around the current pathway. Physical models derived from Maxwell's equation are conventionally used to reconstruct the resistance distribution, which is usually considered to be equivalent to a pressure distribution. Since the number of electrodes required for contact sensing is relatively small, ERT-based tactile sensors have shown many design advantages such as material compatibility [20], [21], fabrication simplicity [16], [18], and resilience to breakage [22].

Despite these benefits, ERT-based tactile sensors can suffer from poor spatial resolution and force estimation accuracy since the reconstruction of the pressure map from voltage measurements is an ill-posed and nonlinear inverse problem [19]. The solution is thus prone to measurement noise and modeling errors, especially related to the electrode contact impedance, electrode misplacement, physical nonlinearity, and inaccurate geometric parameters [14].

To achieve better reconstruction performance compared to the standard physics-based approach, researchers have tried combining electrical simulations of ERT-based sensors with neural networks (NNs). Martin and Choi [23] used an electrical conductivity model to synthesize voltage data from randomized conductivity distributions. Then, they utilized the simulated data to train an artificial neural network (ANN) that effectively learns the nonlinear inverse mapping between the voltage measurements and conductivity distributions. Park *et al.* [24] proposed the EIT-NN framework that adopts a convolutional neural network (CNN) architecture and a specially designed loss function; this framework showed good generalizability for reconstruction by accounting for the spatially varying sensitivity of the ERT sensor. Duan *et al.* [25] used a reconstruction model obtained from an electrical conductivity model combined with spatio-temporal regularization. Randomizing the noise applied to the electrodes improved the spatial resolution of their ERT-based tactile sensor. Although these results are encouraging, it is important to note that these studies all assumed proportionality between the electrical conductivity distribution and the pressure map, which is not generally appropriate.

In order for ERT-based tactile sensors to be useful in real applications, such as robot control and mechanical property characterization, the reconstructed distribution should be expressed in mechanical units, such as force or contact pressure, rather than electrical units, such as resistance or conductivity. This electromechanical calibration process is not trivial because of the nonlinearity of the soft piezoresistive materials that are widely used in ERT-based tactile sensors. Recently, data-driven approaches have been investigated to handle this calibration problem. One straightforward approach is constructing an end-to-end mapping using real-world training data. For instance, Zhang *et al.* [26] trained

a touch-location sensing model with the measurements from experiments wherein human participants touched 16 predefined sensor locations. Similarly, Russo *et al.* [27] classified discrete contact locations using data from single-contact indentation experiments. Although these data-driven approaches can localize a single contact point, their multi-contact localization and contact force estimation have been unsatisfactory. Good performance would require collection of sufficient real data with multiple simultaneous contacts, which is both time consuming and technically challenging.

To overcome the need for copious experimental data in ERT tactile sensors, some researchers have recently begun exploring a sim-to-real transfer approach that combines a physics-based simulation model with data from a small number of experiments so that the mapping can be learned from both simulated and real data. An electromechanical model capable of simulating voltage measurements caused by tactile stimuli can provide observations of a large number of multi-contact cases with lower cost and complexity than physical multi-contact experiments. We previously introduced such a multiphysics model of a standard 20-cm-by-20-cm ERT tactile sensor and applied it to single-contact and double-contact force prediction using transfer learning [28]; the sim-to-real transfer network was trained using only single-contact cases. That approach estimated contact location well on a limited set of 36 possible contact points, but it had poor force prediction accuracy and could not predict more than two simultaneous contacts. These preliminary results showed both the promise and limitations of a simple sim-to-real approach for learning the highly nonlinear and non-localized mapping of an ERT tactile sensor.

This article investigates multi-contact force estimation using a significantly improved modular pipeline that involves two stages of mapping: real-to-sim voltage transfer trained with a moderate number of matched real and simulated multi-contact experiments, and voltage-to-force inverse mapping trained using the simulated voltages calculated from the force data of the real experiments plus a larger amount of more diverse simulation data synthesized from the multiphysics model. For this study, we design a new ERT-based sensor aiming for large-area tactile sensing with simplified manufacturing and electronics compared to our previous designs [16], [18], [28]. Since the multiphysics modeling plays an important role in this framework, we introduce a model consisting of three approximate but sufficiently accurate physics-based models. We thoroughly evaluate our modular pipeline's ability to identify contacts and contact forces using a previously unseen dataset containing both single-contact and multi-contact experiments with the real sensor. Its performance is also quantitatively compared to a pure physics-based mapping approach and a pure end-to-end learning approach to highlight the benefits of combining these two previously developed methods.

## II. EXPERIMENTAL SETUP

Three hardware setups were created for this study: a new ERT-based tactile sensor, an automatic multi-point indenter, and manual indenters. As the implementation details of the

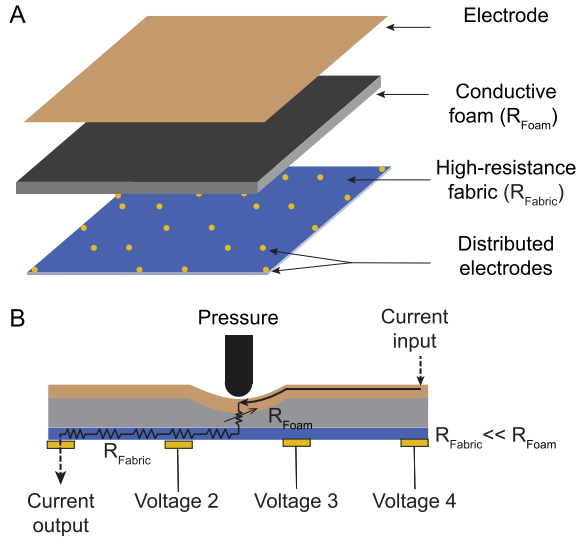


Fig. 1. (A) An overview of the newly developed ERT tactile sensor and (B) its working principle.

tactile sensor are outside the scope of this study, we explain only the essential properties relevant to the multiphysics model of the sensor. Implementation details such as the fabrication process and circuitry can be found in Lee *et al.* [18].

#### A. Tactile Sensor

Our new ERT-based tactile sensor is made of two conductive textile layers surrounding a conductive foam layer (Fig. 1A). It has a rectangular shape and measures 540 mm by 560 mm, which is much larger than previously developed ERT tactile sensors (200 mm by 200 mm in [16], [18], [28]). The sensor's top functional layer is a low-resistance fabric (Medtex P70, Statex, Germany) that has a surface resistivity of  $2 \Omega/\text{sq}$  and functions as a large electrode. The thickness of this fabric is 0.42 mm. The next layer is a conductive open-cell foam (ESD foam, Wolfgang Warmbier, Germany) that has a resistivity of nearly  $10^{11} \Omega/\text{sq}$  when undeformed and a corresponding undeformed thickness of 4 mm; compression closes the air gaps in the foam, bringing its conductive elements into contact and locally reducing the material's resistivity. Under the conductive foam is a sheet of 0.38-mm-thick high-resistance fabric (LTT-SLPA-20K, Eeonyx, USA) with 28 sparsely distributed electrodes arranged in a hexagonal pattern. The distributed electrodes are sewn with conductive thread and have a thickness of nearly 1 mm including the fabric. The resistivity of the high-resistance fabric is  $10^4 \Omega/\text{sq}$ , which is much higher than the low-resistance fabric and much lower than the undeformed foam.

Since the conductive foam has a porous structure, the contact resistance between it and the high-resistance fabric is very large when the sensor is not compressed. When pressure is locally applied to the sensor, the conductive foam is compressed in that region and the local contact resistance at the two interfaces greatly decreases, resulting in a reduction of the resistance between the bottom layer of high-resistance fabric and the top layer of low-resistance fabric (see Fig. 1B). The piezoresistive characteristics of conductive foam structures like

this have been widely investigated due to their low hysteresis and high sensitivity [29].

To perform electrical resistance tomography on the sensor, we inject an electrical current from the top electrode to each of the distributed electrodes in the bottom layer in succession. This current injection scheme always uses the top electrode as a voltage source, similar to the designs of Bera and Nagaraju [30] and Yoshimoto *et al.* [31]. We made the piezoresistive structure using conductive foam to achieve a broader contact-force sensing range, and we placed internal electrodes in addition to boundary electrodes to achieve more uniform sensing performance across the sensor [16]. Compared to our prior ERT-based tactile sensors [16], [18], [28], this new design requires fewer multiplexer circuits and achieves nearly two times higher signal-to-noise ratio in decibels (a factor of ten in magnitude) at the full-scale output (see Suppl. S-I).

The injected current mainly flows through compressed regions of the sensor, if there are any, and forms an electrical potential across the high-resistance fabric; the voltages of all of the distributed electrodes are measured using the adjacent pair protocol [19]. This current injection and voltage measurement process repeats until every distributed electrode has been used once for current injection, which results in a vector of 784 voltage measurements (28 distributed electrodes  $\times$  28 voltage measurements per injection). For this study, we set the sensor's voltage measurement system to run at 24 kHz, yielding an overall update rate of about 30 frames per second. The configuration details of the electronics for this new ERT tactile sensor are provided in Suppl. S-II.

#### B. Multi-Point Indenter

We designed and fabricated a multi-point indenter to conduct indentation experiments and acquire indentation forces along with the corresponding voltage measurements from the sensor. Multi-contact indentation can span an infinite variety of contact location and contact force combinations, so it is physically impossible to cover every combination; instead, we aim to record a reasonable variety of multi-point contact interactions and supplement that real dataset with even more widely varied data from simulation. Since the number of simultaneous contacts tends to be less than four for whole-body humanoids [32], we limited the indenter's maximum number of contact locations to six.

Figure 2A shows the experimental setup containing the ERT tactile sensor, a three-dimensional (3D) stage for Cartesian position control of the indenter, and the multi-point indenter itself. The linear stage precisely controls the position of the center of the multi-point indenter across the surface of the sensor and in the normal direction to cause contact. The multi-point indenter can hold up to six indenter tips that are each instrumented with a load cell (FC22, Measurement Specialties Inc., USA). To vary the relative placement of the contacts, the load cells can move in the radial direction along linear sliders (see Fig. 2B). The radial displacements of the six load cells are controlled together by a single rotary actuator that is linked to a disk with six helical slots. Each load cell has a cylindrical indenter tip with a diameter of 20 mm, which is slightly larger

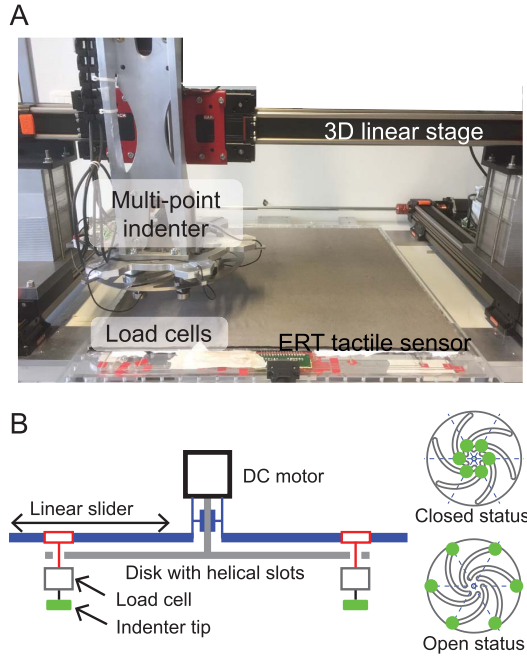


Fig. 2. (A) Experimental setup and (B) schematics of the multi-point indenter.

than the contact area of a human fingerpad. The indenter tips can be manually attached and detached to set the number of simultaneous indentations between one and six contacts.

The multi-point indenter was precisely controlled in a quasi-static manner to avoid dynamic effects. Only normal force was applied, with zero tangential motion after contact initiation. For each indentation case, the contact forces, and sensor voltage measurements were stored together as the data for a single real indentation trial. The average time to conduct each indentation trial was about 3 seconds with the planar motion, radial adjustment, normal indentation, measurement, and retraction.

### C. Manual Indentation System

We additionally prepared manual indenters for conducting multi-contact experiments in random locations, as shown in Fig. 3. An overlay of printed paper marks a hexagonal grid of possible contact locations. This setup allows us to collect a test dataset that is markedly different from the dataset obtained from the multi-point indenter. Every manual indenter has a cylindrical tip with a diameter of 20 mm, identical to the automatic indenter. One 0.230-kg weight and up to three 0.250-kg weights can be loaded over the 0.020-kg tip to make three different weights: 0.500 kg, 0.750 kg, and 1.000 kg, which apply normal forces of 4.905 N, 7.358 N, and 9.810 N, respectively.

## III. MULTIPHYSICS MODEL

The precise description of the mechanical and electrical behavior of a piezoresistive laminate requires complex and computationally demanding models. Building such tools is challenging in terms of modeling and parameter identification,




Fig. 3. The experimental setup for the manual multi-contact experiment; the inset photograph shows one manual indenter with its indenter tip and weights.

but they can then be used to generate simulated data that complement the real data collected from the physical sensor.

Accordingly, we introduce and evaluate an electromechanical multiphysics model of our new ERT-based tactile sensor. Since the sensor has a laminated structure that is subjected to only surface normal pressure, we relate the piezoresistive behavior to the displacement in the normal direction. We also assume that all of the laminate layers are perfectly bonded together. Under these assumptions, the behavior of the sensor can be modeled with three consecutive physics models that connect contact pressure to voltage measurements. The model parameters are either obtained from the literature or calibrated via simple experiments. This multiphysics modeling aims to provide a versatile framework instead of a material-specific one.

### A. Deformation Model

We built a customized finite element (FE) model to calculate the mechanical deformation of the sensor under indentation loads. This methodology allows accurate prediction of the deformations by realistically simulating the propagation and interaction of the internal stresses. Our model is designed to contain an optimal level of detail such that the essential mechanics are captured with only moderate computational cost. Such an approach significantly improves the efficiency with which useful simulated data can be generated.

The FE model of the sensor includes all three material layers shown in Fig. 1. For the computations, we adopted linear elastic solid element formulations and developed in-house FE analysis software in MATLAB (Mathworks, USA) based on the version used in our prior work [28]. Each layer is discretized with 8-node isoparametric brick elements that are assigned isotropic material properties; this treatment is reasonable because the compressive deformation mainly results from the foam. The details of the employed FE formulation can be found in [33].

Instead of generating a typical orthogonal mesh, we construct a tailored mesh with a hexagonal pattern that matches the hexagonal geometry of our multi-point indenter, as shown in Fig. 4A. We predetermined a set of 804 points where the



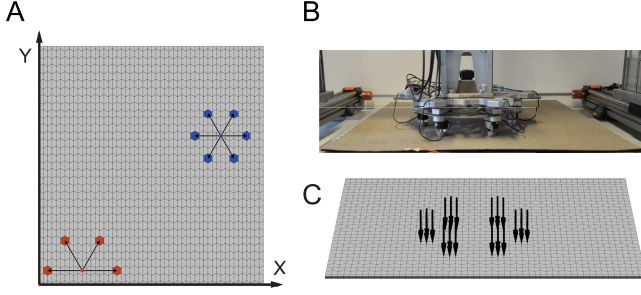


Fig. 4. (A) The hexagonal mesh pattern of 804 predetermined locations where the center of the multi-point indenter and the indenter tips can be located, along with two exemplary indentation configurations involving six and four contacts. (B) A sample indentation experiment involving six contacts on the real sensor and (C) the corresponding finite element deformation model utilized to simulate this experiment.

center of the multi-point indenter and the indenter tips can be located during contact experiments, essentially discretizing the indenter's planar motion and radial expansion. Pre-selecting these indentation points allows us to approximate the contact conditions using point forces and a constant mesh, avoiding the computationally costly and non-robust approach of standard contact analysis involving multiple discretized bodies. The number of indentation points was chosen to obtain sufficiently dense indentation configurations while avoiding high computation time. If these 804 points were arranged in a rectangular grid, it would be approximately 28 by 28, which is far larger than the six-by-six grid of contact points tested in our prior work [28].

Figure 4B shows the experimental setup applying a particular set of six contacts, and Figure 4C shows the corresponding FE simulation. The nodes at the bottom of the sensor are fixed in accordance with the rigid workbench that supports the real setup. The indentation force exerted by an indenter tip in the experimental setup is applied to the model as six point forces at the vertices of the hexagonal mesh at the location of the indenter tip; the measured force is distributed equally to the six selected nodes, as shown in Fig. 4C. This approach accounts for the cylindrical geometry of the indenter tip since the longest diagonal of the hexagons is equal to the indenter diameter.

The material properties used in the FE model are calibrated with two sets of experimentally measured data (see Suppl. S-III). For single-point indentation experiments, the FE model estimates the indented displacement with a normalized root-mean-square error (RMSE) of 24.1% in spite of the nonlinear behavior of the real material (see Suppl. S-IV).

### B. Piezoresistive Layer Model

As illustrated in Fig. 1B, the piezoresistive behavior of the sensor is mainly caused by the compression of the conductive foam (which locally reduces its resistivity) and the local reduction of the contact resistance between the conductive fabric layers and the conductive foam [29]. The contact resistance change is mainly related to interfacial contact area

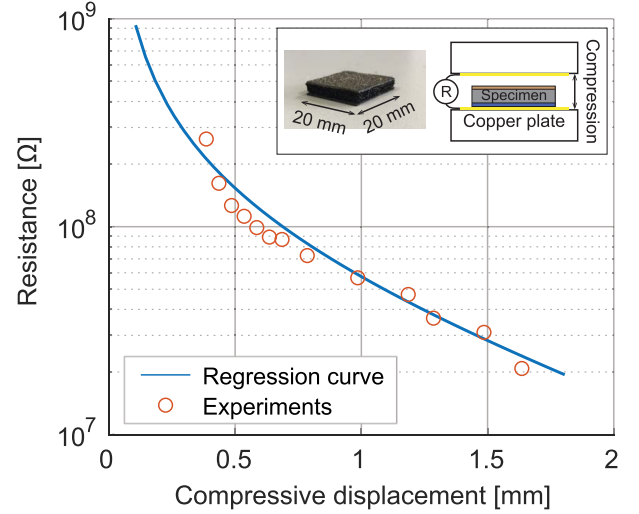


Fig. 5. Experimental results of the compression test and the corresponding regression curve (Eq. 1). Note the log scale on the y axis. A photograph of the specimen and a schematic of the compression experiment appear in the box.

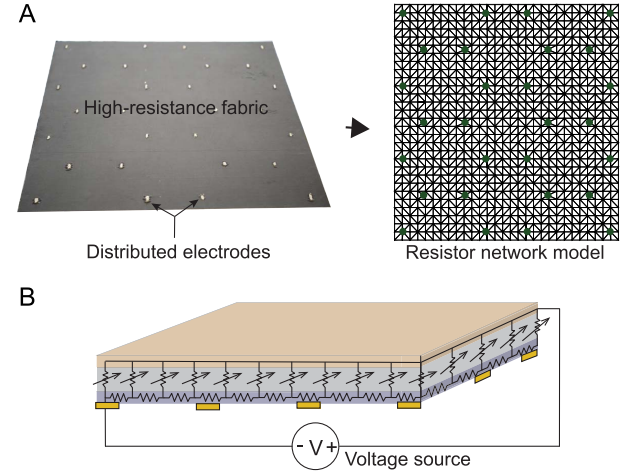


Fig. 6. (A) A photograph of the sensor's high-resistance fabric layer with the 28 distributed electrodes alongside a diagram of the equivalent resistor network model. (B) A schematic of the conductivity model including variable resistors that model the piezoresistive structure of the conductive foam.

and current flow [34]. Since precise physical modeling of the contact resistance change of the elastic foam is nonlinear and highly complex, an approximate empirical model is used instead.

When the conductive foam is not compressed, the contact resistance between the layers is nearly infinite. As the foam is compressed, the resistance rapidly drops and then begins to level off. This behavior can be described as follows:

$$f(x) = \frac{a}{(e^{bx} - 1)} + c \quad (1)$$

where  $x$  is the displacement, with  $x > 0$  mm for compression, and  $f(x)$  is the corresponding contact resistance in  $\Omega$ .

The three coefficients are obtained via regression using experimental data. We prepared a small specimen (20 mm by 20 mm) that represents the working principle of the sensor (see Fig. 5). The specimen was placed between copper electrodes and compressed gradually while we measured its

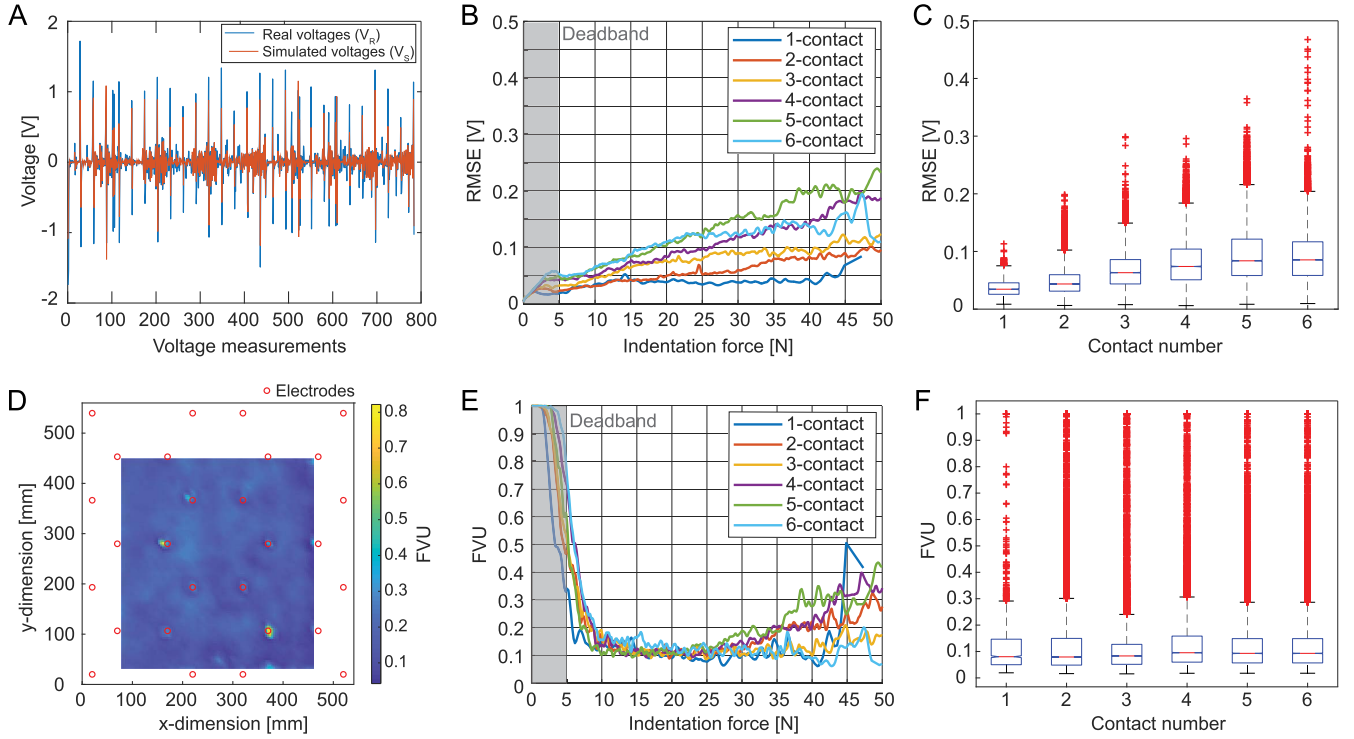


Fig. 7. (A) An example of the 784 real and 784 simulated voltage measurements for a single contact condition. (B) The RMSE between all of the real and simulated voltage measurements as a function of indentation force and the number of contacts. (C) Boxplots of the same RMSE data grouped by contact number. (D) A heat map showing the mean fraction of variance unexplained (FVU) for a rectangular region of the sensor's surface when contacted with only one indenter tip. (E) FVU values as a function of indentation force and the number of contacts. (F) Boxplots of the same FVU data without measurements from the sensor's deadband, organized by contact number.

through-thickness resistance; this experiment thus captures how compression changes both the foam's own resistivity and the contact resistance on its upper and lower surfaces. Figure 5 shows the captured experimental data and the curve fit to the data from Eq. 1 with  $a = 2.2 \times 10^8 \Omega$ ,  $b = 1.4 \text{ mm}^{-1}$ , and  $c = 6.5 \times 10^5 \Omega$ . The curve captures the behavior of the data well; the contact resistance rapidly decreases after initial contact and then begins to level off at higher displacements.

The empirical model converts compressed displacement calculated from the deformation model into the local resistance between the top electrode and the high-resistance fabric; this model element operates as a variable resistor that depends on local deformation. To model piezoresistive behavior across the entire ERT tactile sensor, we assumed an array of parallel variable resistors with these same properties. The top ends of all of these resistors are connected together to represent the low-resistance top electrode, and each bottom end is connected in the appropriate location to a two-dimensional resistor network model of the high-resistance fabric.

### C. Conductivity Model

A conductivity model calculates the electrical potentials that occur across the high-resistance fabric when electrical current is injected from the top electrode layer to one of the distributed electrodes. As shown in Fig. 6A, we first model the high-resistance fabric as a FE mesh including the locations of the distributed electrodes and the array of variable resistors. Then the FE mesh is converted to an equivalent two-dimensional resistor network to simplify inclusion of the

variable resistors [35], [36]. The nodes of the resistor network are the predefined locations of the array of variable resistors.

To simulate the ERT method, a voltage source is connected between the top electrode layer and one of the distributed electrodes (see Fig. 6B). When the sensor is not compressed, the resistances of the variable resistors are extremely high, which results in a nearly constant voltage at all of the distributed electrodes. Once the sensor is compressed, the resistances of the variable resistors rapidly decrease in the region of compression, and a measurable voltage distribution is formed. The voltage of every node is computed from the modified nodal analysis, which is a well-established circuit simulation method [37]. The vector of voltage measurements for this current injection location is saved from the electrode node voltages. For single-point indentation experiments, the conductivity model estimated the voltage distribution of the real sensor with a normalized RMSE of 28.3%, as shown in Suppl. S-IV.

### D. Simulation Accuracy of the Multiphysics Model

The simulation accuracy of the developed multiphysics model is evaluated by comparing real voltage measurements  $V_R$  to simulated voltage measurements  $V_S$ . The real voltages are acquired using the multi-point indenter. The simulated voltages stem from the contact locations and forces measured in those real experiments. An example of the real and simulated voltage measurements is shown in Fig. 7A.

As the first quantitative measure, the RMSE is calculated between the real and simulated voltages from matching contact

TABLE I

OVERVIEW OF THE EXPERIMENTS FOR COLLECTING DATA, WITH THE NUMBER OF TRIALS FOR EACH CONTACT NUMBER RANGING FROM ONE TO SIX, AS WELL AS THE TOTAL NUMBER OF TRIALS. WHETHER REAL OR SIMULATED, EACH EXPERIMENTAL TRIAL YIELDS A LIST OF ONE TO SIX CONTACT LOCATIONS, THE CORRESPONDING LIST OF ONE TO SIX CONTACT FORCES, AND A VECTOR OF 784 VOLTAGES OUTPUT BY THE SENSOR

Experiment	Source	Contact Pattern	Single	Double	Triple	Quadruple	Quintuple	Sextuple	Total
Real-Hex	Multi-Contact Indenter	Hexagonal	20 092	21 313	19 712	14 487	15 944	2 126	93 674
Real-Random	Manual Indenters	Random	0	50	50	50	0	0	150
Sim-Hex	Model	Hexagonal	20 092	21 313	19 712	14 487	15 944	2 126	93 674
Sim-Random	Model	Random	33 275	33 450	33 190	33 047	33 630	33 408	200 000

conditions. Figure 7B shows the average RMSE as a function of the indentation forces (maximum of all contact points); the data is grouped by the number of contacts. The graph shows that the simulation error increases as the indentation force increases from 0 N to 50 N, which makes sense because the voltage signals being predicted generally grow with applied force. Figure 7C shows box plots of the RMSE depending on the contact number. The RMSE increases with the number of contacts up to five contacts.

As the second quantitative measure, the fraction of variance unexplained (FVU) is evaluated between  $V_R$  and  $V_S$ :

$$\text{FVU} = \frac{\sum_{i=1}^K (V_R(i) - V_S(i))^2}{\sum_{i=1}^K (V_R(i) - \bar{V}_R)^2} \quad (2)$$

where  $K = 784$  and  $\bar{V}_R$  is the mean of all of the real voltages measured in a particular experiment. First, the FVU across a portion of the sensor's surface is analyzed from a single-contact experiment to determine how much the behavior of the sensor varies spatially. Figure 7D shows a heatmap of the mean FVU depending on the contact location, annotated with the locations of the distributed electrodes. The FVU tends to be higher near the electrodes than at all other sensing locations, indicating that mechanical contact near some of the electrodes may be modifying the sensor's electrical behavior in that region. Rather than attempting to eliminate this source of spatial variation (which might not be possible), we allow it to persist and expect that our deep neural networks should be able to learn this pattern from the real data on which it is trained.

Figure 7E shows the mean of FVU as a function of the maximum indentation force measured across the six indenter tips. This plot makes it apparent that the ERT tactile sensor has an insensitive range below about 5 N indentation force, so that the FVU is nearly unity in this region; we call this insensitivity to low forces the sensor's deadband. Above 5 N indentation force, the FVU rapidly decreases to about 0.1 and then gradually rises as the indentation force increases. Figure 7F shows box plots of the FVU depending on the contact number, omitting the data from the deadband. The median values of the FVU are about 0.1 regardless of the contact number. Although there are outliers where the simulation fails to predict the real behavior, the model's overall predictive capability is good. Note that this result is evaluated from randomly selected contact locations that include the electrode locations, which

are known to be problematic, and the boundaries of the sensor, which may deviate from the model due to mechanical edge effects.

#### IV. CONTACT DATA ACQUISITION

In this study, we collected data by running two real experiment with the physical setup and two simulated experiments with the multiphysics model, as listed in Table I. Each trial in either experiment type consists of a pressing interaction with the real or simulated sensor, and it yields the  $1 \leq N \leq 6$  two-dimensional locations of contact, the  $N$  forces of contact, and the  $K = 784$  corresponding voltage measurements. As explained in Sec. III-A, only the 804 predetermined contact locations are used in the experiments.

The real data is acquired from the ERT tactile sensor using either the automatic multi-point indenter or the manual indenters. The contact locations of the multi-point indenter are naturally constrained by the indenter's hexagonal design, and the contact forces tend to correlate with one another because all of the indenter tips move together. Consequently, even a very large real dataset collected with the hexagonal indenter would be insufficient to generalize to unconstrained multi-contact situations. The manual indenters are used to collect random multi-contact scenarios only for evaluation.

In contrast to the real dataset, multiphysics simulation can explore diverse contact cases. We conducted two experiments with the model: the first calculates the voltage measurements under all of the contact conditions that occurred in the real experiment. It thus has the same size but took approximately one sixth of the time to acquire, since a contact trial can be simulated faster than it can be physically conducted. The second simulated experiment randomizes both contact locations and forces to create a large volume of data that is more varied than the experimental data collected from the physical setup.

We reconfigure the data collected in these four experiments into the four datasets presented in Table II. The first three (*Real*, *R2S-Transfer*, and *S2S-Recon*) are each designed for a particular learning task. The first is used to train and evaluate the system from end to end on real data. The second captures the gap between the real sensor's voltages and the simulated sensor's voltages. The third dataset targets the mapping between simulated voltages and the associated force map. An additional dataset (*Real-Unseen*) is created from the Real-Random experiment to test whether sim-to-real transfer provides a true advantage for unseen contact situations.

The following two subsections explain the details of our procedures for acquiring real and simulated data.

#### A. Real Data Acquisition From the Experimental Setup

The multi-point indenter collected real experiment data from 3 000 random configurations, each of which included a random location for the center of the indenter and a random radius for the spacing of the indenter tips. The indenter was moved to be just above the sensor surface in each location, and then 5 mm of vertical displacement was applied in increments of 0.5 mm, giving eleven measurements per configuration. In this scheme, 33 000 contact trials were obtained with a single indenter tip setting over the course of approximately 25 hours. The number of attached indenter tips was then manually changed to a new value from one to six, and the same procedure was repeated, giving a total real data collection time of about 150 hours.

Although the indenter tips are identical, their measured indentation forces differ somewhat due to the slightly irregular flatness of the sensor and the indenter. Given the deadband results of Fig. 7B, we removed all contacts for which the indentation force was below 5 N. After this post-processing step, the number of contacts was re-labeled to yield the results of the Real-Hex experiment (Table I), including 93 674 trials.

In addition, 150 trials were conducted using the manual indenters, including 50 double-contact, 50 triple-contact, and 50 quadruple-contact trials. During each contact case, the weights of the manual indenters were randomly changed, and the locations of the manual indenters were randomly chosen from the 804 predefined locations. Figure 3 presents an example of the triple-contact case. We call this experiment Real-Random.

#### B. Simulation Data Synthesis From the Multiphysics Model

To provide a simulated counterpart to the Real-Hex data, we simulated the exact same experimental conditions in the Sim-Hex experiment. Specifically, the simulated sensor's output voltages were computed from the contact locations and contact forces of each of the 93 674 trials in the real dataset. This experiment aids our efforts at real-to-sim transfer learning by providing observations of the constrained real contact scenario in the simulation environment.

Unlike with the real setup, the multiphysics model allows us to apply any set of forces in any locations across the surface of the simulated sensor. Widely varying indentation experiments were thus also simulated by virtually pressing different combinations of the 804 predefined contact points in the Sim-Random experiment. The number of contact points was randomly selected from one to six, the points were selected randomly, and the force magnitude at each point was randomly chosen from the range of 5 N to 40 N. As a result, 200 000 trials were simulated; this process took approximately 23 hours when running in two parallel threads on a CPU (Intel Core i7-8700K@3.7 GHz) with 32 GB RAM (DDR4-2401).

### V. MAPPING FROM VOLTAGE MEASUREMENTS TO MULTI-CONTACT FORCES

We propose a machine-learning pipeline to infer a force map describing multiple contact locations and forces from the

raw sensor voltage measurements in a data-driven manner. For reference, the data automatically collected in the Real-Hex experiment can be used directly to train an end-to-end model that takes real voltages as its input and provides a force map as its output. We define this learning task and the associated dataset *Real* as listed in Table II. However, the indenter design imposes a particular hexagonal pattern on the contact locations, and it causes all of the applied contact forces to increase together. These physical constraints exclude a vast number of other possible multiple-contact patterns. In our experience, the end-to-end trained model would not generalize to other contact patterns; we test this hypothesis in Sec. V-D.

To tackle the problem of data insufficiency, we construct the machine-learning pipeline in a modular two-stage format, as shown in Fig. 8. First, we address the gap between our real experimental setup and our multiphysics model by training a machine-learning model named **Transfer-Net**; it transfers real voltages ( $V_R$ ) to more closely match the simulated voltages ( $V_S$ ) of the multiphysics model for the same contact conditions, essentially learning to undo the non-idealities of our physical setup. The dataset we created to train this real-to-sim network is called *R2S-Transfer* and consists of only the voltages from Real-Hex and Sim-Hex (Table II). Of course, we could also have learned the inverse mapping, from simulation to reality; we chose to learn from real to simulation because this division segments the entire mapping into two pieces we believed would be simpler to learn, and because generating more simulated data is significantly easier than acquiring more real data.

The second step of our pipeline is called **Recon-Net**, which stands for reconstruction network. It learns to convert a set of 784 simulated voltage measurements into the corresponding force map, which is the output of our entire processing pipeline. Each pixel in the force map contains a value representing the total normal force being applied in that region of the sensor. As listed in Table II, the dataset created for this sim-to-sim learning task is called *S2S-Recon*, and it includes the simulated voltages and simulated force maps from both Sim-Hex and Sim-Random. To prepare the pipeline for actual usage with real sensor voltages, we perturb the simulated voltages during training with noise mimicking the errors made by the **Transfer-Net**.

The following subsections detail the two stages of our architecture, describe the combined mapping, evaluate the whole pipeline, and then compare our proposed method to the two common alternative approaches of end-to-end learning and reconstruction using only a physical model.

#### A. Transfer-Net

1) *Network Structure*: The purpose of the Transfer-Net is to model the differences between our real sensor and the simulation model; this inevitable misalignment is sometimes called the sim-to-real gap. The network has a simple structure of multiple fully connected layers that map the 784 real sensor voltages from a single trial to the 784 simulated voltages for the same contact condition. There are five hidden layers that



TABLE II

OVERVIEW OF THE DATASETS, WHICH EACH INCLUDE INPUT DATA AND OUTPUT DATA THAT DEFINE THE ASSOCIATED LEARNING TASK. THESE DATASETS ARE CONSTRUCTED BY COMBINING PORTIONS OF THE EXPERIMENTAL DATA LISTED IN TABLE I.  $V_R$  AND  $V_S$  STAND FOR VOLTAGES FROM THE REAL AND SIMULATED SENSORS, WHILE  $M_R$  AND  $M_S$  INDICATE FORCE MAPS CREATED FROM LISTS OF REAL OR SIMULATED CONTACTS

Dataset Name	Input Variable	Input Experiment(s)	Output Variable	Output Experiment(s)	# of Datapoints
<i>Real</i>	$V_R$	Real-Hex	$M_R$	Real-Hex	93 674
<i>R2S-Transfer</i>	$V_R$	Real-Hex	$V_S$	Sim-Hex	93 674
<i>S2S-Recon</i>	$V_S$	Sim-Hex + Sim-Random	$M_S$	Sim-Hex + Sim-Random	293 674
<i>Real-Unseen</i>	$V_R$	Real-Random	$M_R$	Real-Random	150

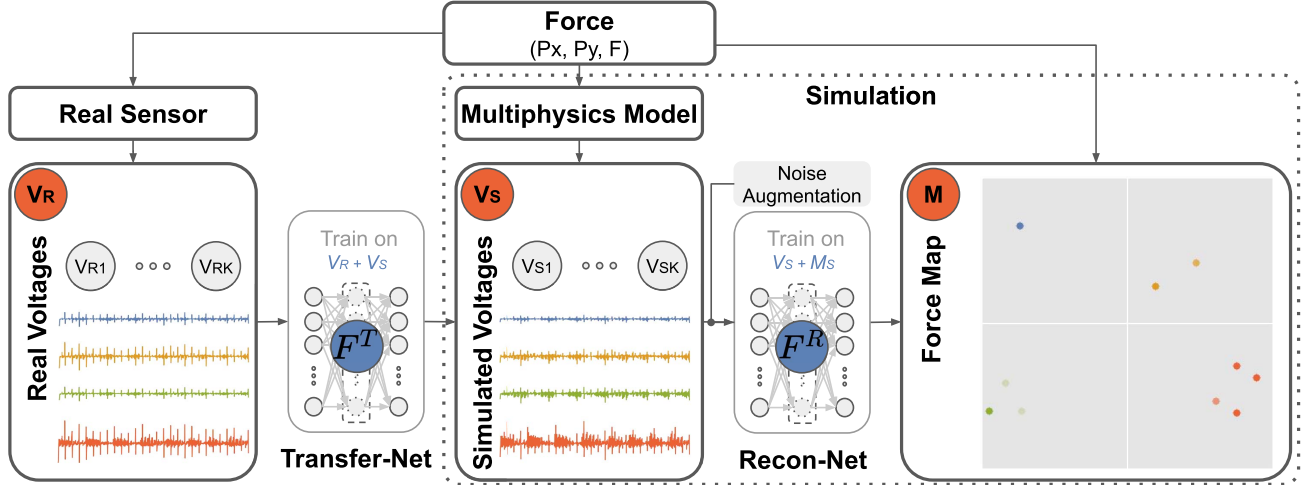


Fig. 8. Two-stage pipeline for mapping real voltages to a force map showing where contact is being made on the surface of the sensor. The first step of this modular pipeline transforms real data measured from our experimental setup so that it closely matches the corresponding data generated by our simulation. The second step operates on the transformed voltages to output the final force map. Four actual sample contact scenarios are shown in blue, yellow, green, and red, to illustrate how the 784 real voltages become 784 voltages similar to the simulation and then yield a force map showing the locations and forces of contact.

each have 784 rectified linear unit (ReLU) neurons. To train the Transfer-Net network, we split the *R2S-Transfer* dataset into training, validation, and test subsets with a ratio of 3:1:1. The Transfer-Net is then trained with the least squares loss function using the training dataset with a batch size of 64 in 1024 epochs. The loss is optimized using Adam with a momentum of 0.7 and a learning rate of 0.001.

2) *Evaluation*: We use RMSE and FVU (Eq. 2) to evaluate the inference performance of the Transfer-Net. As described before, ERT-based sensors have a small number of distributed electrodes and estimate the contacts being applied by measuring the voltage distribution across the piezoresistive sensing surface when current is injected to different locations. Given the mechanical and electrical design of such sensors, the system's voltage measurements during multiple contacts cannot simply be estimated by adding together the voltages measured during the individual contacts; in other words, it is difficult to predict a multi-contact output from only single-contact data.

We validate this challenge of generalization in ERT-based tactile sensors by first training the Transfer-Net using only single-contact data and testing on all multiple contact cases. We then incrementally increase the number of contacts used in the training set and compare the generalization performance of the inference. Table III summarizes these results. As predicted, a network trained on only single-contact data cannot predict

the variance in the voltages that occur with two or more contacts. The prediction performance for unseen multiple-contact cases is greatly improved when trials with that number of contacts are included in the training set. The overall performance is high (less than 6% variance unexplained) for all contact numbers when the training set is representative of the test set. This performance is nearly twice as good as the average FVU values obtained when the Transfer-Net is not applied (Fig. 7E).

### B. Recon-Net

1) *Network Structure*: The Recon-Net is designed to map 784 voltage measurements (output of the Transfer-Net) to the corresponding normal force distribution map. It is trained on simulated data perturbed with naturalistic noise described below. Recon-Net has a deconvolutional network structure that maps from the 784 voltage channels to an image-like force distribution map with  $160 \times 160$  pixels. Each pixel in the map corresponds to a  $4 \text{ mm} \times 4 \text{ mm}$  area of the real sensor describing the distributed force magnitude (in newtons) in the normal direction. Depending on the needed resolution, the size of the force map can be adapted to  $64 \times 64$  or  $32 \times 32$  with increased pixel sizes of 10 mm and 20 mm, respectively.

We train the Recon-Net with the *S2S-Recon* dataset, which contains the data from Sim-Hex (simulated voltages matching the real experiments) and Sim-Random (synthesized force and

TABLE III

RESULTS FOR THE TRANSFER LEARNING TASK (*R2S-Transfer* DATASET). I, II, III, IV, V, AND VI STAND FOR THE PORTIONS OF THE DATASETS WITH THE NAMED NUMBER OF SIMULTANEOUS CONTACTS. GRAY TEXT INDICATES EXTRAPOLATION TO AN UNSEEN NUMBER OF CONTACT POINTS. THE NUMBERS ARE RMSE / FVU ON THE TRAINING SET AND THE DIFFERENT TEST SETS, WITH MEAN RMSE IN VOLTS AND MEAN FVU IN PERCENTAGE

Train	I 0.002 / 1.0						I & II 0.002 / 0.5					
Test	I 0.003 / 4.8	II 0.026 / 75.5	III 0.031 / 64.4	IV 0.047 / 92.8	V 0.055 / 97.8	VI 0.047 / 78.7	I 0.003 / 5.0	II 0.009 / 4.3	III 0.019 / 19.4	IV 0.025 / 20.3	V 0.034 / 34.0	VI 0.029 / 32.1
Train	I, II & III 0.004 / 0.7						I, II, III & IV 0.004 / 0.7					
Test	I 0.004 / 5.1	II 0.009 / 4.7	III 0.011 / 4.1	IV 0.019 / 13.2	V 0.024 / 14.3	VI 0.030 / 16.0	I 0.004 / 5.1	II 0.009 / 4.5	III 0.011 / 4.0	IV 0.010 / 2.3	V 0.019 / 6.2	VI 0.027 / 12.6
Train	I, II, III, IV & V 0.006 / 0.9						I, II, III, IV, V & VI 0.007 / 1.1					
Test	I 0.004 / 5.2	II 0.010 / 4.8	III 0.011 / 4.2	IV 0.010 / 2.4	V 0.012 / 2.1	VI 0.016 / 5.6	I 0.004 / 5.7	II 0.010 / 5.4	III 0.012 / 5.0	IV 0.010 / 2.5	V 0.012 / 2.3	VI 0.014 / 4.2

TABLE IV

RECON-NET EVALUATION ON THE FORCE MAP'S SPATIAL RESOLUTION (PIXEL SIZE) AND FORCE DISTRIBUTION DIAMETER. WE REPORT THE MEAN  $\pm$  STANDARD DEVIATION OF THE SIGNED VALUES FOR FOUR METRICS. NUMBER ERROR IS THE PERCENTAGE BY WHICH THE CONTACT NUMBER IS ESTIMATED INCORRECTLY. LOCATION ERROR IS THE MEAN DISTANCE BETWEEN EACH TRUE CONTACT LOCATION AND THE CLOSEST ESTIMATED CONTACT, WHILE THE FORCE ERROR IS THE MEAN PERCENTAGE BY WHICH THE ESTIMATED FORCE DEVIATES FROM THE GROUND-TRUTH FORCE MEASUREMENT. DIAMETER ERROR IS THE MEAN PERCENTAGE DIFFERENCE BETWEEN THE ESTIMATED CONTACT DIAMETER AND THE GROUND-TRUTH INDENTER DIAMETER

Force Map Pixel [mm] / Diameter [mm]	Map 1 4 / 20	Map 2 4 / 28	Map 3 4 / 32	Map 4 4 / 44	Map 5 4 / 52	Map 6 4 / 60	Map 7 10 / 10	Map 8 10 / 30	Map 9 10 / 50	Map 10 10 / 70	Map 11 20 / 20	Map 12 20 / 60
Number Error [%]	-22 $\pm$ 23	-18 $\pm$ 22	-14 $\pm$ 20	-11 $\pm$ 18	-12 $\pm$ 18	-12 $\pm$ 19	-28 $\pm$ 26	-15 $\pm$ 20	-11 $\pm$ 18	-12 $\pm$ 21	-15 $\pm$ 22	-4 $\pm$ 26
Location Error [mm]	11 $\pm$ 22	11 $\pm$ 18	11 $\pm$ 17	13 $\pm$ 16	13 $\pm$ 14	15 $\pm$ 15	12 $\pm$ 25	14 $\pm$ 22	13 $\pm$ 18	16 $\pm$ 19	15 $\pm$ 26	23 $\pm$ 36
Force Error [%]	-1 $\pm$ 52	-7 $\pm$ 42	-3 $\pm$ 42	5 $\pm$ 43	0 $\pm$ 40	-1 $\pm$ 39	-18 $\pm$ 44	-2 $\pm$ 42	5 $\pm$ 46	-2 $\pm$ 41	-7 $\pm$ 43	-6 $\pm$ 48
Diameter Error [%]	-31 $\pm$ 12	11 $\pm$ 17	36 $\pm$ 18	69 $\pm$ 22	107 $\pm$ 25	134 $\pm$ 28	-34 $\pm$ 21	19 $\pm$ 17	81 $\pm$ 22	171 $\pm$ 28	18 $\pm$ 15	128 $\pm$ 18

voltage pairs). As the contact force information is measured only at the center of each of the  $N$  contact points, we first need to transform these labels to image-like force maps representing the force applied to each small region of the sensor. Two parameters define the transformation: the first is the pixel size, which sets the resolution of the force map, and the second is the diameter of the circular area over which each real indentation force is distributed, to account for the fact that the indenter tips are large relative to the default pixel size. Details about the network structure are provided in Suppl S-V. The training procedure uses the same hyperparameters as the Transfer-Net.

2) *Evaluation*: We evaluate the influence of training data on the performance of the Recon-Net; specifically, we vary the two parameters of the force map (pixel size and force distribution diameter) in twelve different combinations.

The resulting Recon-Nets are evaluated using their prediction errors on four criteria: number of contacts, contact location, contact force, and contact diameter. The input to this evaluation is the produced force map. The number of contacts is counted by finding each pixel whose force magnitude is the highest within its circular neighborhood of 64mm diameter with a 5 N threshold. After extracting the peaks of the estimated contacts in this way, we use the Hungarian matching method to align the predicted contact locations with the ground-truth locations. The force estimation accuracy is evaluated between the matched pairs by comparing the measured contact force with the sum of all of the contact forces predicted in the peak's neighborhood. The contact diameter is calculated from the force map based on the force estimates in the region of each contact using the full width at half maximum (FWHM) criterion, which is similar to RES<sub>50</sub>,

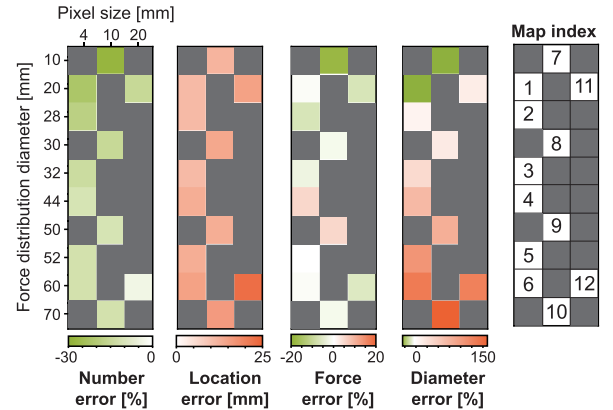


Fig. 9. Visualization of the mean metrics from Table IV to show how pixel size and force distribution diameter affect the performance of the twelve tested Recon-Nets. The gray boxes represent parameter combinations that were not tested.

a widely used metric for ERT tactile sensors [38]; this value is compared to the real indenter tip diameter (20mm).

Table IV shows the performance of the twelve Recon-Nets, and the same results are visualized in Fig. 9. Comparing Map 1 to Map 12, where the pixel size and force distribution size are larger, we see the contact number inference improves slightly, while the contact location and diameter inferences perform worse. Comparing from Map 1 to Map 6, the increase in force distribution size tends to improve the force error by reducing the standard deviations. These two parameters thus allow one to tailor the overall system's performance to emphasize either force measurement (contact number and force) or spatial measurement (contact location and diameter).

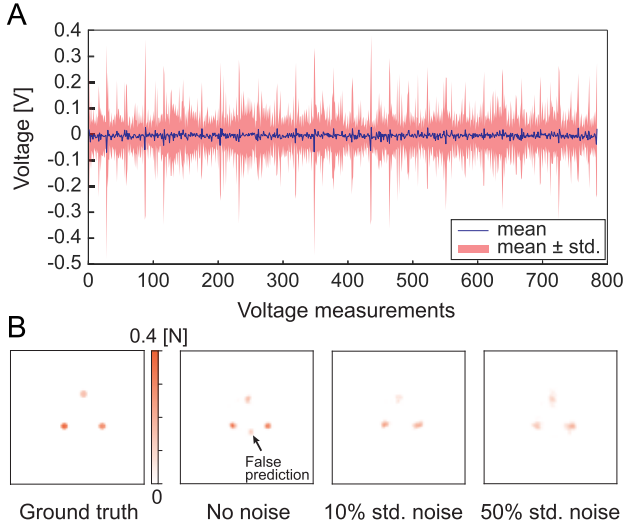


Fig. 10. (A) The mean and standard deviation of the transfer error and (B) force map predictions across the entire surface of the sensor with three different levels of added noise.

### C. Whole-Pipeline Mapping

The whole pipeline combines the Transfer-Net and the Recon-Net in series. If the Recon-Net is trained only on simulated voltages without considering measurement noise, the trained Recon-Net can easily overfit to the simulation data and fail to generalize. Although the Transfer-Net has a low average test error of about 4% FVU on previously unseen real voltage data when it is trained on all contact numbers (Table III), these unpredicted voltage deviations are enough to create highly incorrect force maps because the reconstruction mapping is highly nonlinear and sensitive to small voltage variations.

**1) Reconstruction Robustness for Real-to-Sim Voltage Transfer Error:** To make the reconstruction process robust against the real-to-sim voltage transfer error, we add Gaussian noise to the simulated voltages while the Recon-Net is trained. The standard deviation of each element in the Gaussian noise vector is determined from the residual error of the Transfer-Net, as shown Figure 10A. The mean is close to zero for every one of the 784 measurements, indicating that the training did indeed converge well. However, the standard deviation varies from about 0.05 V for well predicted measurements up to about 0.5 V for a small number of voltage measurements.

The noise magnitude is scaled to control the amount of Gaussian noise added to the simulated voltage measurements in the *S2S-Recon* dataset. Figure 10B showcases a triple-contact situation and its force map predictions trained from three different noise levels. The tested noise levels were 0%, 10%, and 50% of the standard deviations of the transfer error. For this sample, the no-noise scenario predicts a false contact, while the force map of the 10% noise condition is more accurate. The 50% noise prediction is blurred and dilated relative to the 10%. In general, the higher the noise level used during training, the smoother the reconstruction tends to be. The following section evaluates the effect of noise augmentation in detail.

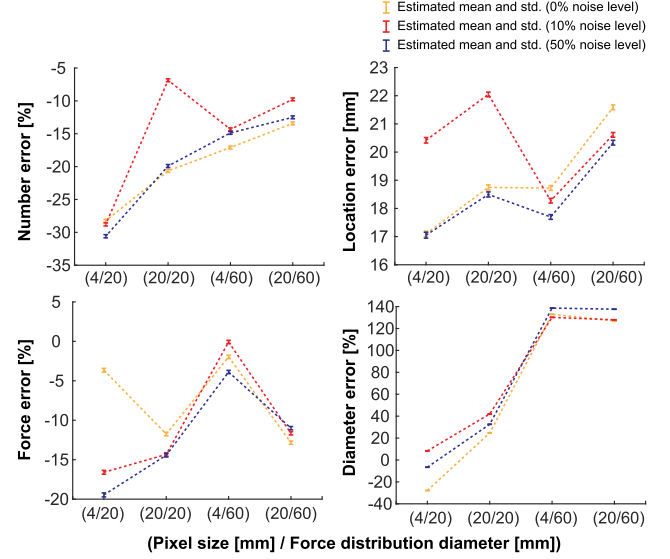


Fig. 11. Force-map inference performance of the whole pipeline for four Recon-Net configurations (4mm and 20mm pixel size, 20mm and 60mm force distribution diameter) with three different noise levels (0%, 10% and 50%).

**2) Evaluation:** The whole pipeline is evaluated using the same four criteria used in the Recon-Net evaluation. The effects of three parameters (force-map pixel size, force-map force distribution diameter, and noise augmentation level) are evaluated with three-factor analysis of variance (ANOVA). Two pixel sizes (4 mm and 20 mm) and two force distribution diameters (20 mm and 60 mm) are chosen from the twelve Recon-Nets for a fully factorial design. The chosen maps are Map 1, Map 6, Map 11, and Map 12. The noise augmentation levels are selected as 0%, 10%, and 50% of the standard deviations of the transfer error. This parameter selection forms a  $2 \times 2 \times 3$  factorial design. For the twelve resulting conditions, we calculate the four performance metrics for the portion of the *Real* dataset that was not used to train the network. Each contact point of a multiple contact case is taken to be an independent sample; thus, the number of samples for each metric is nearly 40000.

The ANOVA evaluated the main effects of the three factors as well as their interactions. Figure 11 shows the estimated means and standard deviations of the four metrics from the ANOVA. For the number error, the main effects of pixel size and force distribution diameter were particularly significant ( $p < 0.0001$ ). Averaged across the three noise levels, the estimated mean of the number error was the smallest for the 20 mm pixel size with 60 mm force distribution diameter ( $-11.9\%$ ); 4 mm pixel size with 20 mm force distribution diameter had the highest average number error ( $-29.2\%$ ). The 10% noise level showed smaller average number error ( $-14.9\%$ ) than the other noise levels ( $-19.8\%$  and  $-19.4\%$  for 0% and 50% noise, respectively).

The main effects of pixel size and force distribution diameter were also particularly significant ( $p < 0.0001$ ) for the remaining three metrics: location error, force error, and diameter error. The estimated mean of the location error was smaller for the 4 mm pixel size (18.2 mm) than for the 20 mm pixel size

(20.3 mm). The 10% noise level showed the largest average location error (20.3 mm compared to 19.0 mm and 18.3 mm for 0% and 50% noise, respectively). The estimated mean of the force error was larger for the 20 mm pixel size with 60 mm force distribution diameter ( $-13.5\%$ ) than the 4 mm pixel size with 60 mm force distribution diameter ( $-1.9\%$ ). The estimated mean of the diameter error was smaller for the 20 mm force distribution diameter (12.1%) than for the 60 mm force distribution diameter (132.3%). The 10% noise level showed the largest diameter error (77.1% compared to 64.2% and 75.6% for 0% and 50% noise levels, respectively).

The noise level showed only a weak effect on the four criteria when the force distribution diameter is 60 mm. 10% noise shows better performance in contact number estimation but worst performance in estimation of contact location and diameter. The estimation of contact force gets worse with the strength of the noise. This performance trend is similar to what happens when the pixel size increases (Table IV).

#### D. Force-Map Prediction for Unseen Multi-Contact Scenarios

We chose to use multiphysics simulation in tandem with deep learning to try to overcome the data-deficiency problem for random multi-contact cases that are impractical to acquire in the real world. The *Real-Unseen* dataset is used to evaluate the whole pipeline's performance for unseen contact situations. For this evaluation, we also prepared two comparison models that are widely used for force prediction: direct end-to-end learning from the *Real* dataset and physics-based reconstruction (see Suppl. S-VI for further details on both comparison methods).

Figure 12A shows two example cases of the force-map predictions obtained from our hybrid approach, end-to-end learning, and physics-based reconstruction. Our real-to-sim method predicts the stronger contacts seen in the ground-truth labels. In contrast, the end-to-end learning approach shows a complete mismatch between predictions and ground truth, although it used the same Recon-Net structure and is trained on the same real data as our approach. The physics-based approach tends to show a blurry force map with blobs roughly centered around the ground-truth contacts with very low predictions for contact force magnitude; its qualitative performance is marginally coherent with the ground truth.

Figure 12B shows the performance of our real-to-sim approach and end-to-end learning evaluated with three of our four metrics: contact location error, contact force error, and contact diameter error. The contact number error is not plotted because the location error of the end-to-end approach is high. The physics-based reconstruction is not included in this comparison because its blurry outputs result in erratic evaluation metrics and high errors. Note that the results shown for both our approach and end-to-end learning were obtained from the twelve different Recon-Net configurations introduced in Sec. V-B; we show all three noise levels for real-to-sim.

The results demonstrate that the force-map prediction performance of our pipeline is highly superior to that of both baseline approaches; the only metric on which end-to-end

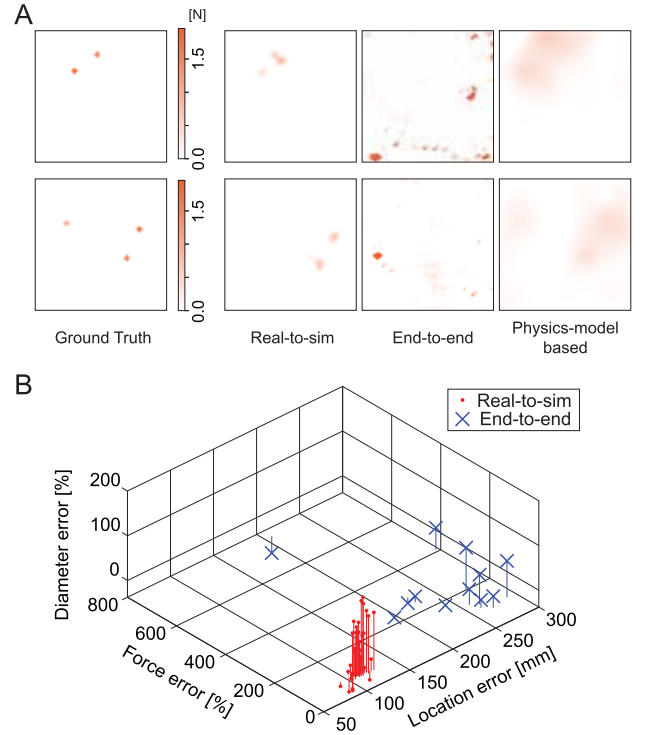


Fig. 12. (A) Two example cases of the force-map prediction performance using three different approaches (real-to-sim, end-to-end, and physics-based model). (B) Performance comparison between the real-to-sim obtained from the 36 Recon-Nets (Map 1 to Map 12 with three noise levels) and the end-to-end approaches obtained from the twelve Recon-Nets (Map 1 to Map 12). The physics-based reconstruction is excluded from this comparison due to its inaccurate results.

learning marginally beats our method is the contact diameter prediction error. Among the different tested variations of our transfer-learning approach, we define the best performance to be the minimum average location error and close to zero average force and diameter error. According to this standard, the Recon-Net with Map 1 and 50% noise augmentation showed the best performance: it achieves average errors of  $-58.3\%$  in contact number, 74.4 mm in contact location (about 13.5% of the 560 mm sensor width), 17.2% in contact force, and  $-9.3\%$  in contact diameter.

## VI. DISCUSSION

ERT tactile sensors are a promising technology for realizing a practical tactile skin; however, estimating the force map from such a sensor's voltage measurements has been a challenging problem due to its nonlinear and ill-posed characteristics. This study introduces an ERT tactile sensor and investigates a way to obtain an accurate multi-contact force-mapping network by combining multiphysics simulation and real experiments.

#### A. ERT Tactile Sensor Design

Common ERT tactile sensors switch both the current source electrode and the current sink electrode [14], [16]. The ERT tactile sensor presented in this study includes one large electrode that is always used as a current source. This current injection strategy was chosen to gain practical advantages, such as simplifying manufacturing and reducing measurement

electronics. The resulting tactile sensor was shown to function well, achieving a signal-to-noise ratio almost double that of the common design. Future studies should compare these two designs more thoroughly, including the mathematical properties of each one's inverse voltage-to-force mapping.

A sensor design that provides strong superposition with locality in the voltage-to-force map would simplify transfer learning because a multi-contact case could then be generalized as a combination of single contacts. Unfortunately, existing ERT tactile sensors do not have this property, such that patterns learned from single contacts do not generalize well to more contacts, as listed in Table III. Nevertheless, addition of the multi-contact dataset seems to enable some extrapolation to cases with more contacts. This result indicates that a sufficient number of real measurements are necessary for an ERT tactile sensor to achieve satisfactory transfer performance even when the contacts are constrained to follow a given spatial pattern. This limitation may be resolved by adding more electrodes to the ERT sensor or optimizing its measurement pattern [17].

### B. Multiphysics Model

Among the three physics-based components of the multiphysics model, the piezoresistive model is the most important for capturing the nonlinear behavior of the sensor. The deformation model and the conductivity model are relatively easy to calibrate because they are linear. Although the experimental result in Fig. 5 presented the piezoresistive characteristics of the sensor, its exact characteristics can depend on the sensing location due to the stochastic nature of the contact between the conductive foam and the high-resistance fabric; we expect Transfer-Net learns to reduce these variations in the real voltages.

In terms of efficiency, the real data acquisition took 2.73 seconds per trial, while the current multiphysics model took 0.41 seconds per trial, which is nearly six times faster. It should be noted that this relatively fast computation time was possible only because we assumed the sensor experienced only normal indentation, without shear. This assumption enabled us to simplify the complex piezoresistive behavior of the multi-layer, multi-material sensor structure. To extend this concept to complex deformation cases, a more sophisticated multiphysics model may be required, which could greatly increase the computation time. In that case, reduced models could be considered for deformation and conductivity [39].

### C. Training Data Acquisition

Regarding the experimental setup, the six indenter tips on the multi-point indenter are controlled in three-dimensional space by only four actuators. This design imposes strong geometrical constraints on the data collected in the Real-Hex experiment. This situation is similar to a case where an autonomous robot learns the tactile perception capabilities of its own skin through sensitive fingertips. The role of the real-to-sim transfer module is to overcome the geometrical constraints of the real experimental setup, so that the final system can correctly perceive widely varying contact patterns. A limitation of the end-to-end learning approach is

clearly seen in Fig. 12; since the ERT tactile sensor has low generalization capabilities from single-contact to multi-contact scenarios, end-to-end direct learning trained on the geometrically constrained *Real* dataset fails to predict unseen random multi-contacts.

### D. Machine-Learning Pipeline

Compared to the previous approaches that used conductivity simulation with machine-learning techniques [23]–[25], the proposed pipeline transfers real voltages to be similar to simulated voltages and then produces a reconstructed force map as its output. The relationship between conductivity and force is nonlinear, as shown in Fig. 5; therefore, conductivity simulation approaches cannot output force information without additional non-linear processing by neural networks. Our method produces a force map that can represent multiple contacts of different sizes and shapes; this approach is more sophisticated and potentially more useful than previous efforts that localized only a single contact point using classifiers [27].

We introduced three parameters of the whole-pipeline structure that can influence the multi-contact force map inference performance: force-map pixel size, force distribution diameter, and noise augmentation level. These parameters seem to regularize the ill-posed Recon-Net. The results shown in Fig. 9 indicate that increasing the pixel size and the force distribution diameter increases the error of contact location and contact diameter estimates. However, force estimation error and contact number prediction error showed the opposite relationship. Increasing the pixel size reduces the resolution of the force map. For instance, the 4 mm pixel size discretizes the sensor into a force map of  $160 \times 160$ , so Recon-Net has an output size of 25 600. Training a mapping from 784 voltage measurements to 25 600 pixels of the force map is an under-determined problem. Increasing the pixel size helps mitigate this problem of under-determination. Regarding the force distribution diameter, this parameter should be selected to be sufficiently large such that every pixel of the force map experiences non-zero force values during the training. Increasing the force distribution diameter can relieve the sparsity problem, although predictions of contact position and diameter become worse.

As seen in Fig. 11, adding a small amount of noise (10%) to the simulated voltages generally helped the force-map prediction performance, while a large amount of noise (50%) was less helpful; in contrast, adding noise generally increased the force estimation error. It is known that augmenting the training data with noise is equivalent to a regularization process [40]. In our approach, the real-to-sim voltage transfer error is modeled as Gaussian noise. If the characteristics of the transfer error were more carefully investigated, another noise model might be able to improve the overall performance.

### E. Prediction for Unseen Multi-Contact Scenarios

The results of the unseen random multi-contact test clearly demonstrate that the proposed real-to-sim transfer-learning approach achieves substantially better force-map prediction



performance on unseen multi-contact cases than both baselines. The physics-model-based baseline was obtained by linearizing the physics model, and the end-to-end learning baseline was trained only on real data. Although iterative approaches are known to produce better performance for nonlinear mappings [19], we did not consider such approaches due to their heavy computation time, which is not suitable for ERT tactile sensors with high bandwidth. We expect the results of our approach would be even better if the final performance evaluation had spanned the sensor's entire range from 5 N to 40 N rather than only low forces between 4.905 N and 9.810 N. Omitting the paper sheets that were used to locate the manual indenters might also improve the predictions.

## VII. CONCLUSION

Estimating a multi-contact force map from voltage measurements is a challenging problem in ERT-based tactile sensors whether one uses data-driven approaches or physics-based models; the mapping is nonlinear, and the reconstruction process is ill-posed. This paper introduced a hybrid approach that learns a data-driven mapping from real voltages of a new soft ERT tactile sensor and simulated voltages synthesized from a multiphysics model. The low-cost multiphysics model simulates voltage measurements caused by tactile stimuli with only a 10% fraction of variance unexplained (FVU), even for multi-contact cases. The data-driven mapping pipeline is made of two stages: a real-to-sim voltage transfer network trained with a moderate number of real multi-contact experiments subject to geometric constraints plus a voltage-to-force reconstruction network trained using a larger amount of randomly simulated voltages augmented with noise. For validation, the pipeline is evaluated using a previously unseen dataset containing multi-contact experiments conducted with the real sensor. When tested with these 150 unseen multi-contact cases, the proposed pipeline showed an average error of 74.4 mm in contact position (which is about 13% of the sensor width), 17.2% force estimation error, and 9.3% contact diameter prediction error. This result outperforms both a conventional physics-based mapping approach and direct end-to-end learning. The approach presented in this study has the potential to greatly ease the development process for many robotic tactile skins by simulating a broad range of tactile inputs that are not subject to the constraints or costs of real experiments.

## DATA AVAILABILITY STATEMENT

The raw data collected in this study and the Python code for the presented machine-learning framework are available from Edmond, the open research data repository of the Max Planck Society (<https://dx.doi.org/10.17617/3.8p>).

## ACKNOWLEDGMENT

The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Huanbo Sun.


## REFERENCES

- [1] R. A. Russell and S. Parkinson, "Sensing surface shape by touch," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 1993, pp. 423–428.
- [2] A. Schmitz, P. Maiolino, M. Maggiali, L. Natale, G. Cannata, and G. Metta, "Methods and technologies for the implementation of large-scale robot tactile sensors," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 389–400, Jun. 2011.
- [3] P. Mittendorf and G. Cheng, "Humanoid multimodal tactile-sensing modules," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 401–410, Jun. 2011.
- [4] D. Anghinolfi, G. Cannata, F. Mastrogianni, C. Nattero, and M. Paolucci, "On the problem of the automated design of large-scale robot skin," *IEEE Trans. Autom. Sci. Eng.*, vol. 10, no. 4, pp. 1087–1100, Oct. 2013.
- [5] C. Bartolozzi, L. Natale, F. Nori, and G. Metta, "Robots with a sense of touch," *Nature Mater.*, vol. 15, no. 9, p. 921, Aug. 2016.
- [6] L. Van Duong and V. A. Ho, "Large-scale vision-based tactile sensing for robot links: Design, modeling, and evaluation," *IEEE Trans. Robot.*, vol. 37, no. 2, pp. 390–403, Apr. 2021.
- [7] M. Ramuz, B. C.-K. Tee, J. B.-H. Tok, and Z. Bao, "Transparent, optical, pressure-sensitive artificial skin for large-area stretchable electronics," *Adv. Mater.*, vol. 24, no. 24, pp. 3223–3227, Jun. 2012.
- [8] B. Ward-Cherrier *et al.*, "The TacTip family: Soft optical tactile sensors with 3D-printed biomimetic morphologies," *Soft Robot.*, vol. 5, no. 2, pp. 216–227, Apr. 2018.
- [9] G. Vásárhelyi, B. Fodor, and T. Roska, "Tactile sensing-processing: Interface-cover geometry and the inverse-elastic problem," *Sens. Actuators A, Phys.*, vol. 140, no. 1, pp. 8–18, Oct. 2007.
- [10] N. F. Lepora, U. Martinez-Hernandez, M. Evans, L. Natale, G. Metta, and T. J. Prescott, "Tactile superresolution and biomimetic hyperacuity," *IEEE Trans. Robot.*, vol. 31, no. 3, pp. 605–618, 2015.
- [11] K. Flores De Jesus, M. H. Cheng, L. Jiang, and E. G. Bakhoun, "Resolution enhancement method used for force sensing resistor array," *J. Sensors*, vol. 2015, no. 647427, Jan. 2015.
- [12] H. Sun and G. Martius, "Machine learning for haptics: Inferring multi-contact stimulation from sparse sensor configuration," *Frontiers Neurobotics*, vol. 13, p. 51, Jul. 2019.
- [13] A. Yao and M. Soleimani, "A pressure mapping imaging device based on electrical impedance tomography of conductive fabrics," *Sensor Rev.*, vol. 32, no. 4, pp. 310–317, 2012.
- [14] D. Silvera-Tawil, D. Rye, M. Soleimani, and M. Velonaki, "Electrical impedance tomography for artificial sensitive robotic skin: A review," *IEEE Sensors J.*, vol. 15, no. 4, pp. 2001–2016, Apr. 2015.
- [15] K. Liu *et al.*, "Artificial sensitive skin for robotics based on electrical impedance tomography," *Adv. Intell. Syst.*, vol. 2, no. 4, Apr. 2020, Art. no. 1900161.
- [16] H. Lee, K. Park, J. Kim, and K. J. Kuchenbecker, "Internal array electrodes improve the spatial resolution of soft tactile sensors based on electrical resistance tomography," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, May 2019, pp. 5411–5417.
- [17] K. Park, H. Lee, K. Kuchenbecker, and J. Kim, "Adaptive optimal measurement algorithm for ERT-based large-area tactile sensors," *IEEE/ASME Trans. Mechatronics*, vol. 27, no. 1, pp. 304–314, Feb. 2022.
- [18] H. Lee, K. Park, J. Kim, and K. J. Kuchenbecker, "Piezoresistive textile layer and distributed electrode structure for soft whole-body tactile skin," *Smart Mater. Struct.*, vol. 30, no. 8, Aug. 2021, Art. no. 085036. [Online]. Available: <http://iopscience.iop.org/article/10.1088/1361-665X/ac0c2e>
- [19] D. S. Holder, *Electrical Impedance Tomography: Methods, History and Applications*. Boca Raton, FL, USA: CRC Press, 2004.
- [20] F. Visentin, P. Fiorini, and K. Suzuki, "A deformable smart skin for continuous sensing based on electrical impedance tomography," *Sensors*, vol. 16, no. 11, p. 1928, 2016.
- [21] H. Lee, D. Kwon, H. Cho, I. Park, and J. Kim, "Soft nanocomposite based multi-point, multi-directional strain mapping sensor using anisotropic electrical impedance tomography," *Sci. Rep.*, vol. 7, no. 1, pp. 1–10, Feb. 2017.
- [22] H. Lee, K. Park, Y. Kim, and J. Kim, "Durable and repairable soft tactile skin for physical human robot interaction," in *Proc. Int. Conf. Hum.-Robot Interact. (HRI)*, 2017, pp. 183–184.
- [23] S. Martin and C. T. M. Choi, "Nonlinear electrical impedance tomography reconstruction using artificial neural networks and particle swarm optimization," *IEEE Trans. Magn.*, vol. 52, no. 3, pp. 1–4, Mar. 2016.

- [24] H. Park, K. Park, S. Mo, and J. Kim, "Deep neural network based electrical impedance tomographic sensing methodology for large-area robotic tactile sensing," *IEEE Trans. Robot.*, vol. 37, no. 5, pp. 1570–1583, Oct. 2021.
- [25] X. Duan, S. Taurand, and M. Soleimani, "Artificial skin through super-sensing method and electrical impedance data from conductive fabric with aid of deep learning," *Sci. Rep.*, vol. 9, no. 1, pp. 1–11, 2019.
- [26] Y. Zhang, G. Laput, and C. Harrison, "Electrick: Low-cost touch sensing using electric field tomography," in *Proc. CHI Conf. Hum. Factors Comput. Syst.*, May 2017, pp. 1–14.
- [27] S. Russo, R. Assaf, N. Carbonaro, and A. Tognetti, "Touch position detection in electrical tomography tactile sensors through quadratic classifier," *IEEE Sensors J.*, vol. 19, no. 2, pp. 474–483, Jan. 2019.
- [28] H. Lee, H. Park, G. Serhat, H. Sun, and K. J. Kuchenbecker, "Calibrating a soft ert-based tactile sensor with a multiphysics model and sim-to-real transfer learning," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 1632–1638.
- [29] M. Panahi-Sarmad *et al.*, "A comprehensive review on carbon-based polymer nanocomposite foams as electromagnetic interference shields and piezoresistive sensors," *ACS Appl. Electron. Mater.*, vol. 2, no. 8, pp. 2318–2350, Aug. 2020.
- [30] T. K. Bera and J. Nagaraju, "Common ground method of current injection in electrical impedance tomography," in *Proc. Global Trends Inf. Syst. Softw. Appl.*, 2012, pp. 574–587.
- [31] S. Yoshimoto, Y. Kuroda, and O. Oshiro, "Tomographic approach for universal tactile imaging with electromechanically coupled conductors," *IEEE Trans. Ind. Electron.*, vol. 67, no. 1, pp. 627–636, Jan. 2020.
- [32] J. Borràs, C. Mandery, and T. Asfour, "A whole-body support pose taxonomy for multi-contact humanoid robot motions," *Sci. Robot.*, vol. 2, no. 13, Dec. 2017, Art. no. eaaq0560.
- [33] G. R. Liu and S. S. Quek, *The Finite Element Method: A Practical Course*, vol. 1. Oxford, U.K.: Butterworth-Heinemann, 2003.
- [34] S. J. Proctor, L. W. Linholm, and J. A. Mazer, "Direct measurements of interfacial contact resistance, end contact resistance, and interfacial contact layer uniformity," *IEEE Trans. Electron Devices*, vol. ED-30, no. 11, pp. 1535–1542, Nov. 1983.
- [35] L. Borcea, V. Druskin, and F. G. Vazquez, "Electrical impedance tomography with resistor networks," *Inverse Problems*, vol. 24, no. 3, Jun. 2008, Art. no. 035013.
- [36] M. Capllonch-Juan and F. Sepulveda, "Evaluation of a resistor network for solving electrical problems on ohmic media," in *Proc. 11th Comput. Sci. Electron. Eng. (CEECE)*, Sep. 2019, pp. 35–40.
- [37] F. N. Najm and R. C. Dumas, *Circuit Simulation*, vol. 9. Hoboken, NJ, USA: Wiley, 2010.
- [38] D. S. Tawil, D. Rye, and M. Velonaki, "Improved image reconstruction for an EIT-based sensitive skin with multiple internal electrodes," *IEEE Trans. Robot.*, vol. 27, no. 3, pp. 425–435, Jun. 2011.
- [39] B. Peherstorfer and K. Willcox, "Data-driven operator inference for nonintrusive projection-based model reduction," *Comput. Methods Appl. Mech. Eng.*, vol. 306, pp. 196–215, Jul. 2016.
- [40] C. M. Bishop, "Training with noise is equivalent to Tikhonov regularization," *Neural Comput.*, vol. 7, no. 1, pp. 108–116, 1995.



**Hyosang Lee** received the B.S. degree in mechanical engineering from Korea University, Seoul, Republic of Korea, in 2010, and the M.S. and Ph.D. degrees in mechanical engineering from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea, in 2012 and 2017, respectively. He is currently a Research Scientist at the Max Planck Institute for Intelligent Systems, Stuttgart, Germany. His research interests include tactile sensors and physical human–robot interaction.



**Huanbo Sun** received the B.S. degree in mechanical engineering from Southeast University Chengxian College, China, in 2013, and the M.S. degree in mechanical engineering from RWTH Aachen University, Germany, in 2016. He is currently pursuing the Ph.D. degree with the Max Planck Institute for Intelligent Systems, Tübingen, Germany. His research interests in designing machine-learning-driven haptic sensors include hardware, software, and theory development.



**Hyunkyu Park** received the B.S. and M.S. degrees in mechanical engineering from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea, in 2017 and 2019, respectively, where he is currently pursuing the Ph.D. degree. His research interests include soft robotics, deep learning, and wearable robotics.



**Gokhan Serhat** (Member, IEEE) received the B.S. degree in mechanical engineering from Middle East Technical University, Ankara, Turkey, in 2011, the M.S. degree in computational mechanics from the Technical University of Munich, Germany, in 2014, and the Ph.D. degree in mechanical engineering from Koc University, Istanbul, Turkey, in 2018. He is currently a Research Scientist at the Max Planck Institute for Intelligent Systems, Stuttgart, Germany. His research interests include computational mechanics, numerical methods, and mechanical vibrations.



**Bernard Javot** (Member, IEEE) received the Ph.D. degree in haptics from Université Pierre et Marie Curie (UPMC), Paris, France, in 2016. He had before ten years of experience as a Mechanical Engineering Teacher in Paris (agrégé de génie mécanique). He is currently a Research Engineer at the Haptic Intelligence Department, Max Planck Institute for Intelligent Systems, Stuttgart, Germany. His expertise includes actuator design and sensor integration.



**Georg Martius** received the degree in computer science from the University of Leipzig in 2005 and the Ph.D. degree from the University of Göttingen, Germany, in 2009. He was a Post-Doctoral Fellow at the IST Austria and a Post-Doctoral Researcher at the Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany. He has been leading the research group on autonomous learning at the Max Planck Institute for Intelligent Systems, Tübingen, Germany, since 2017. His research is focused on machine learning for robotics, including internal

model learning, reinforcement learning, intrinsic motivations, representation learning, differentiable combinatorial optimization, and haptics.



**Katherine J. Kuchenbecker** (Fellow, IEEE) received the Ph.D. degree in mechanical engineering from Stanford University in 2006. She did her post-doctoral research at JHU. She was a Faculty Member at the GRASP Laboratory, University of Pennsylvania, from 2007 to 2016. She is currently the Director at the Max Planck Institute for Intelligent Systems, Stuttgart, Germany. Her research blends robotics and human–computer interaction with foci in haptics, teleoperation, physical human–robot interaction, tactile sensing, and medical applications. She

has delivered a TEDYouth talk on haptics in 2012 and has been honored with the 2009 NSF CAREER Award, the 2012 IEEE RAS Academic Early Career Award, the 2014 Penn Lindback Award for Distinguished Teaching, and various best paper, poster, demonstration, and reviewer awards. She has co-chaired the IEEE RAS Technical Committee on Haptics from 2015 to 2017 and the IEEE Haptics Symposium in 2016 and 2018.