



Learning human–environment interactions using conformal tactile textiles

Yiyue Luo^{1,2,3}, Yunzhu Li^{1,3} , Pratyusha Sharma^{1,3}, Wan Shou^{1,3} , Kui Wu^{1,3}, Michael Foshey^{1,3}, Beichen Li^{1,3}, Tomás Palacios^{2,3}, Antonio Torralba^{1,3} and Wojciech Matusik^{1,3}

Recording, modelling and understanding tactile interactions is important in the study of human behaviour and in the development of applications in healthcare and robotics. However, such studies remain challenging because existing wearable sensory interfaces are limited in terms of performance, flexibility, scalability and cost. Here, we report a textile-based tactile learning platform that can be used to record, monitor and learn human–environment interactions. The tactile textiles are created via digital machine knitting of inexpensive piezoresistive fibres, and can conform to arbitrary three-dimensional geometries. To ensure that our system is robust against variations in individual sensors, we use machine learning techniques for sensing correction and calibration. Using the platform, we capture diverse human–environment interactions (more than a million tactile frames) and show that the artificial-intelligence-powered sensing textiles can classify humans' sitting poses, motions and other interactions with the environment. We also show that the platform can recover dynamic whole-body poses, reveal environmental spatial information and discover biomechanical signatures.

Living organisms extract information and learn from the surroundings through constant physical interactions¹. Humans are particularly receptive to tactile cues (on hands, limbs and torso), which allow complex tasks such as dexterous grasp and locomotion to be carried out². Observing and modelling interactions between humans and the physical world are thus fundamental to the study of human behaviour³, and also for the development of applications in healthcare⁴, robotics^{5,6} and human–computer interactions⁷. However, studies of human–environment interactions typically rely on easily observable visual or audible datasets⁸, and obtaining tactile data in a scalable manner remains difficult.

Wearable electronics have benefited from innovations in advanced materials^{9–11}, designs^{12–16} and manufacturing techniques^{17,18}. However, sensory interfaces that offer conformal, full-body coverage and can record and analyse whole-body interactions have not been developed. Such full-sized tactile sensing garments could be used to equip humanoid robots with electronic skin for physical human–robot collaboration^{2,6}, could serve as auxiliary training devices for athletes by providing real-time interactive feedback and recording¹⁹, and could assist high-risk individuals, such as the elderly, in emergencies (a sudden fall, for example) and early disease detection (heart attacks or Parkinson's disease, for example)²⁰ by acting as unobtrusive health-monitoring systems.

In this Article, we report a full-body tactile textile for the study of human activities (Fig. 1). Our textiles are based on coaxial piezoresistive fibres (conductive stainless-steel threads coated with a piezoresistive nanocomposite), produced using an automated coating technique. The functional fibres can be turned into large-scale sensing textiles that can conform to arbitrary three-dimensional (3D) geometries, using digital machine knitting, an approach that addresses current challenges in the large-scale manufacturing of functional wearables. An artificial intelligence (AI)-based computational workflow is also developed here to assist in the calibration,

recording and analysis of full-body human–environment interactions using the tactile textiles.

Conformal tactile sensing textiles

Our low-cost (~US\$0.2 per metre; details about cost estimation are provided in Supplementary Table 1) coaxial piezoresistive fibres are created using a scalable automated fabrication process (Fig. 1a; further details on characterization and preparation are provided in Fig. 2a–c and the Methods). A pair of fibres are then orthogonally overlapped to create a sensing unit, which converts pressure stimuli (a normal force acting on the surface) into electrical signals^{1,8}. Figure 2d shows that the measured resistance of a typical sensor drops from ~8 k Ω to 2 k Ω in response to an increasing applied normal force (or pressure) in the range of 0.1–2 N (with an average peak hysteresis of ~22%; Supplementary Fig. 1f). Our functional fibre has reliable performance over 1,000 load and unload cycles (Fig. 2f). It also maintains stable performance in a wide range of daily environments, such as different temperatures (20–40 °C) and humidities (40–60%, Supplementary Fig. 1i–k). The performance of the sensing unit can be tuned on demand by adjusting the material compositions (copper and graphite weight percentages, Supplementary Fig. 1c,d) and fabrication process (pulling speed and material feeding rate, Supplementary Fig. 1e).

To fabricate 3D conformal tactile textiles in a scalable manner, we use digital machine knitting to seamlessly integrate the functional fibres into shaped fabrics and full-sized garments. Although weaving has been widely attempted for inserting functional fibres into fabric, knitting is chosen here as it has two primary advantages over weaving. First, the fabrication process is simpler: whereas woven fabric must be cut and sewn to form a garment, full-garment machine knitting can directly manufacture wearables with arbitrary 3D geometry²¹. Second, the interlocking loops of yarn (that is, stitches) used in knitting create a softer, stretchier fabric, which ensures comfort and compatibility during natural human motions (Supplementary Video 1).

¹Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA. ²Microsystems Technology Laboratories, Massachusetts Institute of Technology, Cambridge, MA, USA. ³Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology, Cambridge, MA, USA. ✉e-mail: liyunzhu@mit.edu; wanshou@mit.edu; torralba@csail.mit.edu; wojciech@csail.mit.edu




Fig. 1 | Textile-based tactile learning platform. **a**, Schematic of the scalable manufacturing of tactile sensing textiles using a customized coaxial piezoresistive fibre fabrication system and digital machine knitting. A commercial conductive stainless-steel thread is coated with a piezoresistive nanocomposite (composed of polydimethylsiloxane (PDMS) elastomer as the matrix and graphite/copper nanoparticles as the conductive filler). **b**, Digitally designed and automatically knitted full-sized tactile sensing wearables: (i) artificial robot skin, (ii) vest, (iii) sock and (iv) glove. **c**, Examples of tactile frames collected during human-environment interactions and their applications explored using machine learning techniques.

Because of its relative stiffness compared to regular knitting yarn (Supplementary Fig. 1g), the piezoresistive functional fibre is not suitable to be knitted directly by forming loops. We thus used an alternative knitting technique—inlay—which horizontally integrates the yarn in a straight configuration (Supplementary Fig. 5e and Supplementary Video 2). We assembled two knitted fabrics with functional fibres inlaid perpendicularly to form a sensing matrix. Different knitting patterns can be used to tune the performance of the sensing matrix for customized applications. The highest achieved sensitivity of 1.75 kPa and a detection range up to 87.5 kPa are demonstrated in Fig. 2e (the configuration is shown in Supplementary Fig. 5f).

We computationally designed and automatically knitted several example garments: socks (216 sensors over an area of 144 cm², Supplementary Fig. 6a), a vest (1,024 sensors over 2,100 cm², Supplementary Fig. 6b), a robot arm sleeve (630 sensors over 720 cm², Supplementary Fig. 6c), and full-sized gloves (722 sensors over 160 cm², Supplementary Fig. 6d). Thanks to the digital machine knitting system²², these garments are fully customizable: they can be adapted to individual shape, size, surface texture and colour preference, meeting the needs of personalization and fashion design. Details of knitting operation and our designs can be found in Supplementary Figs. 5 and 6.

Self-supervised sensing correction

During the large-scale manufacturing of wearable sensors, it is inevitable that variations will be introduced and thus failure.

Although researchers conventionally attempt to fabricate flawless sensor arrays^{23,24}, we draw inspiration from living organisms, which develop the ability to adapt their sensory system in the presence of localized defects or environmental variations using overall sensing^{1,2,4}. We adopt a similar mechanism to provide robust sensing capabilities while relaxing the flawless requirement in sensor fabrication. This is critical for many applications, as it is impractical to perform individual calibration and correction of our sensing units due to their high-density, complex geometries and diverse applications. To implement such a mechanism, we developed a self-supervised learning paradigm that learns from weak supervision, using spatial-temporal contextual information to normalize sensing responses, compensate for variation and fix malfunctioning sensors. In particular, we first calibrate our tactile socks and gloves by collecting synchronized tactile responses and readings from a digital scale stepped or pressed by a wearer (Fig. 3a and Supplementary Video 3). At each frame, the scale reading indicates the overall applied pressure, which is expected to correlate linearly with the sum of tactile responses at all sensing points. We then train a fully convolutional neural network with four hidden layers, which takes a small sequence of raw tactile array responses as input and outputs a single frame with the same spatial array resolution (Supplementary Fig. 8a). The output represents the corrected tactile response of the middle frame of the input sequence. The neural network is optimized via stochastic gradient descent with the objective consisting of two components: one encourages the output to preserve the spatial




Fig. 2 | Characterization of the functional fibre. **a**, Photograph of the piezoresistive functional fibre (>100 m) and sensing fabrics. **b**, Optical image of a stainless-steel thread, coaxial piezoresistive fibre and acrylic knitting yarn. **c**, Scanning electron microscopy (SEM) cross-sectional image of the piezoresistive fibre. **d**, The resistance profile of a typical sensor (composed of two piezoresistive fibres) in response to pressure (or normal force). Error bars indicate the standard deviation of four individual sensors. **e**, The influence of fabric structures ('manual' inlay and automatic inlay) on device performance. Behaving as a buffer, the introduction of the soft fabric (red curve) decreases the sensitivity and increases the sensing range. Also, the ribbed structures obtained from automatic inlay induce gaps between two aligned fabrics (dark blue and light blue curves), further decreasing the sensitivity and increasing the detection range. Detailed structures are shown in Supplementary Fig. 5f. **f**, The stable performance of a typical sensor over 1,000 cycles of force load and unload. (Additional characterization of the fibre and fabric is provided in the Supplementary Information).

details from the input, and the other requires the summed tactile response to be close to the reading from the scale. Our network consistently increases the correlation between the tactile response and the reference (reading from scale)—the correlation improves from 75.9% to 91.1% in the right sock, from 92.4% to 95.8% in the left sock and from 77.7% to 88.3% in the glove (Fig. 3a and Supplementary Fig. 9a–d).

To further calibrate the sensing fabrics with arbitrary shapes, such as the vest and robot arm sleeve, we treat the corrected glove as a mobile, flexible 'scale' and record data from both the tactile glove and the target garments while researchers randomly press on

the latter through the glove. With the same self-supervised learning framework (Fig. 3b), the correlation between the responses increases from 32.1% to 74.2% for the tactile vest and from 58.3% to 90.6% for the tactile robot arm sleeve (Fig. 3b and Supplementary Fig. 9e). The self-supervised calibration network exploits the inductive bias underlying the convolutional layers²⁵, learns to remove artefacts, and produces more uniform and continuous responses (Fig. 3c–h and Supplementary Fig. 10). It enables the large-scale sensing matrix to be robust against variation among the individual elements and even their occasional disruption, consequently improving the reliability of the measurement.




Fig. 3 | Self-supervised correction. **a**, Procedure for correcting the tactile sock. The wearer steps on a digital scale (a tennis ball was placed between the foot and the scale to enhance conformal contact) and synchronized readings from the sock and scale are collected. The same method is used for tactile gloves. **b**, Procedure for self-supervised correction of the vest using the response of the calibrated glove as the reference. The same method is also used for the robot arm sleeve. The bar plots on the right (**a,b**) show the correlations between the tactile response and the scale reading: 'Raw' indicates the correction of the original, unprocessed tactile signal; 'Manual' indicates the correlation of manual-adjusted data where all saturated tactile signals were clipped; and 'Self-supervised' indicates the correlation obtained after self-supervised correction. **c-h**, Examples of raw and corrected readouts from the sock (**c,d**), vest (**e,f**) and robot arm sleeve (**g,h**). Our method removes artefacts and enhances the smoothness of the sensors. The colour bar indicates the relative pressure in each sensing point.

Learning on human–environment interactions

Using our full-body sensing garments, we collected a large tactile dataset (over 1,000,000 frames recorded at 14 Hz, details are provided in the Methods)²⁶ featuring various human–environment interactions, including diverse contacting patterns from sitting on chairs of different materials with different postures, complex body movement and other daily activities (Supplementary Videos 3–7). The capture of human behaviour, skills and crafts is essential for cultural preservation, transfer of knowledge, as well as for human and robot performance optimization^{6,19}. We demonstrate the capability and utility of our platform by leveraging our data for environment and action classification, motion pattern discovery and full-body human pose prediction.

Our full-sized sensing vest shows the characteristic pressure distributions during sitting, standing, reclining and other actions, which indicate the wearer's pose, activity and the texture of the contacted surfaces (Supplementary Video 6). We captured a dataset for a wearer performing different poses over various surfaces (Fig. 4a and Supplementary Fig. 11a). Projecting the high-dimensional sensory responses into 2D space via *t*-distributed stochastic neighbour embedding (*t*-SNE)²⁷, we observe that the recordings from different classes naturally form distinctive clusters, indicating the discriminative power of the vest (Fig. 4b). Our vest also exhibits a discriminative resolution of 2 cm, which is superior to a human's back (~4.25 cm)²⁸. We dressed a manikin in the vest and collect a dataset by pressing three letters cutouts—'M', 'T' and 'T'—against the back




Fig. 4 | Learning on human-environment interactions. **a**, Example photographs and tactile frames. **b**, *t*-SNE plot from our pose dataset recorded by the tactile vest. The separation of clusters corresponding to each pose illustrates the discriminative capability of the sensing vest. **c**, Example photographs and tactile frames of 'M', 'I' and 'T' pressed on the tactile vest. **d**, Letter classification accuracy drops as sensor resolution decreases (top). The confusion matrix for classifying the letter and the orientation (bottom). **e**, Location of the 19 joint angles representing body pose in our model. **f**, Mean squared error in pose prediction. **g,h**, Influence of sensor resolution (**g**) and the number of input frames (**h**, temporal window) on prediction performance. The dashed lines represent the baseline defined as the canonical mean pose of the training data. **i**, Comparison of various poses recorded by MOCAP (ground truth) and the same poses recovered by our model from the tactile socks pressure frames. Notable errors in the arm region are highlighted in red. **j**, Time series prediction of walking. **k**, PCA on tactile maps from walking (insets are corresponding tactile frames).

in different orientations (Fig. 4c and Supplementary Fig. 11c). The classification network (Supplementary Fig. 8b) takes a small window of tactile responses and predicts the type and orientation of the letter with an accuracy of 63.76%; the accuracy drops as the effective resolution decreases (Fig. 4d). The discriminative capability of the sensing socks was demonstrated similarly with action classification. Details are provided in Supplementary Fig. 11e.

Humans maintain the dynamic balance of the body by redirecting the centre of mass and exerting forces on the ground, which results in distinct pressure distributions on the feet^{3,29,30}. Given this, we hypothesize that a person's pose can be estimated from the change of pressure distribution over time obtained by our tactile socks, shown as a sequence of pressure maps (Fig. 4i,j). Here, the body pose is represented by 19 joint angles spanning over the legs, torso and arms (Fig. 4e). We record synchronized tactile data from a pair of sensing socks and a full-body motion capture (MOCAP) suit, while the user conducts versatile actions (Supplementary Video 4). We model the pose prediction task as a regression problem using a convolutional neural network (Supplementary Fig. 8c). The model processes a time-series of tactile array footprints that contain the evolving information about the contact events and predicts the human pose in the middle frame. We optimize the neural network by minimizing the mean squared error between the predicted and the ground truth joint angles (MOCAP data) using stochastic gradient descent.

The model learns to make accurate predictions that are both smooth and consistent over time, achieving a mean squared error that is 70.1% lower than our baseline, which always outputs the mean pose. As shown in Fig. 4f–j and Supplementary Video 5, our model achieves higher accuracy for the poses in the torso and legs than in the arms. This is congruous with our observation that the pressure distributions on the feet are mostly affected by lower body movement and the majority of body mass is located in the torso³. We also observe that the recovered dynamic motion implies the environmental spatial information (Supplementary Video 5), for example, whether the stairs are upwards or downwards. The significance of sensing resolution and temporal information is reiterated, as the performance drops with a systematic reduction in either the input resolution of the tactile pressure map or the context size of the input tactile frames (Fig. 4g,h).

To further understand the patterns that emerged from the tactile dataset recorded by the sensing socks, we used principal component analysis (PCA) to project the tactile signals collected from walking into a 2D plane (Fig. 4k). Intriguingly, the walking signals naturally form a circular pattern in the reduced PCA space. The pressure distribution smoothly transitions back and forth between the left and right foot as we traverse the circle, which describes signatures of the different phases during walking³.

Our results indicate that a pair of tactile socks can potentially replace the bulky MOCAP system. Our approach sets a path towards the analysis and study of human motion activities without much physical obstruction in domains like sports, entertainment, manufacturing activities and care of the elderly. Furthermore, such footprints contain dynamic body balance strategies, demonstrating a valuable, instructive paradigm for robot locomotion and manipulation^{3,6}.

Finally, we also demonstrate that our 3D conformal sensing textiles can be used as tactile robot skin. Most modern robots rely solely on vision; however, large-scale and real-time tactile feedback is critical for dexterous manipulation and interaction skills, especially when vision is occluded or disabled^{31,32}. Our functional textile enables conformal tactile coverage of the robot arm (Fig. 1b–i), gripper, limbs and other functional parts with complex 3D geometries, endowing the robots with real-time tactile feedback (Fig. 3e and Supplementary Video 7). Our platform offers a critical ingredient for unobtrusive multi-point collision detection and physical human–robot interaction, which remains challenging with embedded torque sensors in the robot arm⁶.

Conclusions

We have reported large-scale tactile sensing textiles with dense sensor arrays that are conformal to arbitrary 3D geometries. The textiles are manufactured via an automated, inexpensive and scalable approach that combines functional fibres and full-garment digital machine knitting. A self-supervised sensing correction has been developed to normalize the sensor responses and correct malfunctioning sensors in the array. Demonstrations of various human–environment interaction learning using the system suggest that our approach could be of use in biomechanics, cognitive sciences and child development, as well as in imitation learning for intelligent robots.

Methods

Fabrication of coaxial piezoresistive functional fibre. The functional fibre was constructed in two parts: a conductive core and a piezoresistive sheath. The piezoresistive sheath was prepared by mixing graphite nanoparticle (400–1,200 nm, US Research Nanomaterials), copper nanoparticle (580 nm, US Research Nanomaterials) and PDMS elastomer (Sylgard-184, base to curing agent weight ratio of 10:1, Dow Corning). Silicone solvent OS2 (Dow Corning) was added to optimize the viscosity of the mixture for the coating. The mixture was thoroughly mixed by a speed mixer (FlackTek) at 2,500 r.p.m. for 90 s. The prepared piezoresistive composite was loaded into a syringe (Nordson), which was connected to a customized material reservoir with 500- μ m-diameter inlet and 700- μ m-diameter outlet (Fig. 1a). Constant pressure (20 psi) was applied to the syringe with a Nordson dispenser, while the 3-ply stainless-steel thread (Sparkfun, DEV-13814) was fed to the inlet (Fig. 1a). The thread was then pulled by a continuously rotating motor and coated with the nanocomposites. After thermal curing at 150 °C, the resulting coaxial piezoresistive fibre (namely, the functional fibre) was wound into a roll (Fig. 1a). Each sensor was constructed at the intersection of two orthogonally overlapped functional fibres, forming a layered structure with the piezoresistive nanocomposite sandwiched by two conductive electrodes (inset, Fig. 2d).

Morphological characterization. The longitudinal uniformity is illustrated by optical images (Olympus SZ61, Supplementary Fig. 1a). The microstructure of the coaxial piezoresistive fibre was further characterized using a scanning electron microscope (Zeiss Merlin, Supplementary Fig. 1b). All fibres were characterized without additional conductive coating in the scanning electron microscope using a voltage of 3–5 kV.

Electrical characterization. The individual sensor, constructed from two orthogonally aligned functional fibres, was used for electrical characterization. The resistance profile was recorded by a digital multimeter (DMM 4050, Tektronix) while an adjustable pressure (or normal force) was applied to the sensor by a mechanical testing system (Instron 5944). The applied load was controlled at a specific strain rate for all tests (0.05–0.5 mm min⁻¹). To evaluate the robustness of the sensor, we cycled the force between 0.1 and 0.5 N for more than 1,000 cycles at a constant strain rate of 0.1 mm min⁻¹ (Fig. 2f). We measured the performance of sensors composed of fibres with various graphite and copper compositions and coating thicknesses (Supplementary Fig. 1a,b,h). Sensor stability was studied with different fibre alignment angles, humidity and temperature (Supplementary Fig. 1i–k). The effect of fabric structure on single sensor performance was evaluated by recording the resistance of an isolated sensor embedded in knitted fabrics with different structures. Also, sensors consisting of different combinations of fabric structures inlaid with functional fibres (automatic and manual inlaying) were investigated (Fig. 2e and Supplementary Fig. 5f). The typical sensing resolution ranged from 0.25 cm² to 4 cm² (for details see Supplementary section 1.4).

Mechanical characterization. The tensile test was conducted on an Instron mechanical tester (Instron 5984). The fabricated functional fibre (diameter of 600 μ m) was compared with stainless-steel core thread and two common kinds of acrylic knitting yarn (Tamm Petit C4240 and Rebel TIT8000), as shown in Supplementary Fig. 1g. Samples with a length of 10 cm were pulled at a strain rate of 5 mm min⁻¹. The yield strength of the functional fibre is over six times larger than that of the acrylic knitting yarns. However, its ultimate strain is less than 10% of acrylic knitting yarn's ultimate strain. These mechanical characteristics require special accommodations during the machine knitting process. The functional fibre shows mechanical properties similar to that of the core stainless-steel thread. More mechanical characterizations of the tactile sensing fabric are provided in Supplementary section 1.4.

Conformal 3D sensing textiles manufacturing. *Digital machine knitting.* The full-sized conformal garments and coverings were computationally designed using KnitPaint²² and the stitch meshing framework, and later automatically fabricated with a V-bed digital knitting machine (SWG091N2, Shima Seiki), which has two

beds of needles (front and back) forming an inverted 'V' shape. Each needle is composed of a hook, which catches the yarn and holds the topmost loop, and a slider, which can be actuated to move vertically to close the loop and assist in holding the loop. Movement of the yarn carrier is synchronized to the needle operation, where yarns are fed in with appropriate tension (Supplementary Fig. 5a). We used three basic needle operations during fabrication: knit, tuck and transfer. A knit operation actuates the needles to grab the fed yarn from the yarn carrier, forms a new loop, and pulls it through the existing loop to connect the loops in columns (Supplementary Fig. 5b). A tuck operation actuates the needles to grab the yarn and hold it without forming a new loop (Supplementary Fig. 5c). A transfer operation actuates needles from both beds to pass the existing loop from one bed to the other (Supplementary Fig. 5d). We also utilized racking, where the back bed shifts laterally to the left or to the right as a whole to create needle offsets during transferring to guide the complex 2D and 3D shaping.

Inlaying. Two methods of inlaying were utilized to seamlessly integrate the piezoresistive fibres into conformal 3D knitted textiles: automatic inlaying and manual inlaying. Automatic inlaying requires the fabric structure of 'ribs', which are textured vertical stripes created by alternating columns of knit stitches on the front and back bed. The functional fibre is forced to move simultaneously with normal knitting yarn, which is caught by the needles on the two beds and forms alternating knit stitches to hold down the inlaid functional fibre (Supplementary Fig. 5e and Supplementary Video 2). The ribbed structure (alternating stitches on the front and back bed) allows the straightforward inlaying design; however, it creates a textured gap when two fabrics are aligned orthogonally to act as a functioning device and lowers the sensitivity and accuracy. Manual inlaying requires consecutive movements of the normal knitting yarn and the functional fibre. Four steps of operations are conducted: (1) knitting with normal knitting yarn; (2) the transfer of specific loops from front to back bed; (3) moving the functional fibre across the fabric; (4) transfer of specific loops from back to front bed to their original positions (Supplementary Fig. 5e and Supplementary Video 2). A flat inlaid fabric (without ribbed structure) was fabricated with manual inlay, and designs with continuous functional fibre coverage were achieved by alternating the transfer direction (from the front bed to back bed or from back bed to front bed). However, due to the fibre stiffness and the limited space between two beds, the functional fibre can barely stay down during the second round of stitches transfer and is easily caught by the needles, leading to the destruction of fibre functionality. Therefore, the manual inlay can only be applied to 2D structures.

Full-sized conformal 3D sensing textiles. Two knitted fabrics of specific 2D and 3D shapes with horizontally and vertically inlaid functional fibres were arranged as a double-layer structure to form the large-scale sensing matrix (Supplementary Fig. 6). To optimize the manufacturability and sensing performance of the garments, we exploited both automatic inlaying and manual inlaying. Different combinations of fabric texture (ribbed or not ribbed) were selected for different devices to obtain the desired sensitivity and detection range (Fig. 2e). Full-sized socks, vest, gloves and conformal robot arm (LBR iiwa, KUKA) sleeve were fabricated with the designs shown in Supplementary Fig. 6. To connect the embedded sensing matrices to the readout circuit for data transmission without disrupting the main knitted fabric, cutting was performed at the edge of the fabric, where 1–2-cm functional fibres were pre-reserved through additional tucking at the edge for connection. A modified electrical-grounding-based circuit architecture²⁶ was used to eliminate most crosstalk and parasitic effects of the passive matrix (Supplementary Fig. 7a).

Self-supervised sensing correction. Dataset. During data collection for the tactile sock correction, our researcher wore the sock (with a customized printed circuit board mounted on the calves) and conducted random stepping on a digital weighing scale (Etekcity). The tactile learning platform is powered by a portable energy source (Dell XPS 13 9380, 30 Wh), which allows real-time recording (data acquisition and serialization) while being carried by the researcher in a backpack. The fully charged laptop allows 4 h of recording. We recorded both the real-time tactile information from the sensing matrix and the readings from the scale. To ensure the conformal contact of all sensors, we placed a tennis ball on the scale during the data-collection process (Fig. 3a and Supplementary Video 3). Sensing correction for the tactile gloves used a similar data-collection process. For data collection with gloves, the printed circuit boards were mounted on the arms. A person wore the glove and conducted random pressing on a digital scale (Lansheng). We used 3D printed shapes (dots, lines and curves) to ensure conformal contact of individual sensors (Supplementary Video 3). In total, we collected 25,024 frames for the left sock correction, 37,275 frames for the right sock correction and 108,305 paired frames for glove calibration. We split the dataset sequentially, with the first 80% reserved for training, another 10% for validation and the remaining 10% for testing.

Data processing and network architecture. The data-collection process is assumed to be quasistatic, so the readings from the scale should have a linear correlation with the sum of all sensor responses on the glove or the sock. Based on this observation, we developed a self-supervised learning framework that uses the scale's reading as weak supervision to guide correction of the garment. We preprocessed all frames

by replacing shorted or dead sensors with an average value of its surrounding. A fully convolutional neural network augmented with skip connections was developed (Supplementary Fig. 8a). The model takes a small sequence of the processed frames as input and then adds an individual bias term to each sensor. The output of the model (denoted conv4 in Supplementary Fig. 8a) is a tensor of the same spatial resolution as the input with a channel size of 16. The scale prediction is then defined as the mean of all values in this tensor. We applied an affine transformation to the first channel of the tensor to derive the calibration result, which corresponded to the middle frame of the input sequence. The affine transformation was parameterized by two scalar numbers—a scaling factor and a bias term—which were shared by all sensors in the same garment.

The calibration network was expected to predict the readings from the scale accurately and preserve the details of the input frame. Hence, the overall objective was a reweighted sum of two components: one was the mean squared error between the predicted scale reading and the actual scale reading, while the other was an L1 distance between the calibration result and the corresponding input frame, where we scaled the second reconstruction loss with a factor of 0.1. We optimized the parameters in the convolutional layers, the position-wise bias and the global affine transformation using stochastic gradient descent, which employed the Adam optimizer with an initial learning rate of 0.001 and a batch size of 64. We decreased the learning rate with a factor of 0.8 whenever the loss on the validation set plateaued for five epochs.

Self-correction. After properly correcting the glove, we used it as our new 'scale' to calibrate other garments like the vest and the robot arm sleeve. During data collection, the customized printed circuit board was mounted on the leg for the vest and mounted on the robot arm for the sleeve. We randomly pressed the target garment embedded with the sensing matrix using the corrected glove (Supplementary Video 3). In total, 56,031 paired frames and 44,172 frames were collected for the vest and robot arm sleeve correction. We used the same data split, network architecture and training procedure, except that we used the sum of the sensors from the calibrated glove as the target value to minimize the scale prediction error.

Human pose classification. To evaluate the discriminative power of the vest, and the quality of the obtained signal, a neural-network-based classifier was developed to distinguish different sitting postures and contacted surfaces. As shown in Supplementary Fig. 11a, we selected 10 different classes of poses and recorded 82,836 tactile frames, in total, to construct the dataset. The input to the model is a small sequence of 45 consecutive frames, which totals ~3 s in real time. The sequences were selected from the dataset with a stride of 2, among which 5,000 sequences were used for validation (500 per class), another 10,000 sequences were reserved as the test set (1,000 per class), and we used the rest for training. As shown in Supplementary Fig. 8b, we passed each tactile frame through three shared convolutional layers. The resulting 45 vectors with a length of 512 after flattening were then passed through a bidirectional gated recurrent unit³³ and two fully connected layers to derive the final probabilistic distribution over the 10 classes. We trained the model using the standard cross-entropy loss for 20 epochs with the Adam optimizer with learning rate of 0.001 and batch size of 32. We obtained an accuracy of 99.66% on the test set in distinguishing different lying postures and supporting surfaces (Supplementary Fig. 11b). A real-time pose classification with an accuracy of 94% is demonstrated in Supplementary Video 8. Here, the real-time classifier predicts the probability of each class, accompanied by the ground-truth human pose.

Letter classification. The dataset was collected from the tactile vest worn by a manikin when the models of three letters—"M", "T" and "T"—were pressed against the back (Supplementary Fig. 11c). We recorded 62,932 frames in total for the 10 classes and retrieved a top-1 accuracy of 63.76% and a top-3 accuracy of 88.63% using the same network architecture, data splitting and training procedure as the pose classification task (Supplementary Fig. 11d). To ablate the resolution of the sensor, we reassigned the value in each 2 × 2 grid with the average of the four values, which reduced the effective resolution to 16 × 16. We then used the same classification training pipeline to obtain accuracy. A similar procedure was employed for calculating the accuracy for sensors with effective resolutions of 8 × 8, 4 × 4 and 1 × 1.

Action classification and signature discovery. In this experiment, we used tactile information from the two socks worn by a person to identify which action the wearer was conducting. The dataset consists of tactile frames retrieved from a pair of socks when the wearer conducts nine different activities: walking, climbing up the stairs, climbing down the stairs, fast walking, standing on toes, jumping, leaning on the left foot, leaning on the right foot and standing upright (Supplementary Fig. 11e). The dataset contains 90,295 frames across the different action categories. The same classification pipeline achieved a top-1 accuracy of 89.61% and a top-3 accuracy of 96.97% (Supplementary Fig. 11f).

To further analyse the patterns underlying the signals collected from the socks, we performed PCA to visualize their distribution. We used the 12,245 frames from the class of walking in the action classification dataset. We concatenated

and flattened the signals from the left and right sock as the high-dimensional representation at each time step, each of which has 2,048 dimensions. We extracted the principal components with the highest and the second-highest variance to project the high-dimensional responses to a 2D space and visualize them in Fig. 4k.

Full-body pose prediction. *Dataset.* The key idea of full-body pose prediction is to look at how different sequences of tactile footprints correlate to the different poses as the person transitions from one pose to another. To this end, we simultaneously collected data from a pair of tactile socks (frame rate of 13–14 Hz) and an XSENS motion capture system (MOCAP, frame rate of 50 Hz) worn by a person. The person conducted exercises and other daily activities, including walking, bending forward, twisting, lunging and so on (Supplementary Video 4). The collected dataset is diverse in terms of poses as the person transitions through different tasks. The MOCAP is composed of 17 inertia-based sensors, mounted on 17 key points on the human body to record and estimate the orientation of 19 different joints during the movement. The real-time pressure imprints from both feet were recorded and fed into the network. The 19 different joints include the joints of the legs, arms and torso. The collected dataset includes 282,747 frames of concurrent MOCAP and tactile pressure maps: 236,036 frames (~83.5%) were used as the training set, 10,108 frames (~3.5%) as the validation set and 36,603 frames (~13%) as the test set.

Data processing and network architecture. We trained a deep convolutional neural network to predict the 19 different joints of the human body, given the tactile footprints of the person. The architecture of the network is described in Supplementary Fig. 8c. The network consists of two convolution layers, which take in a sequence of tactile frames from the socks from time step $t-k$ to time step $t+k$. The input to the model is 30 consecutive frames of the pressure map of the left and right feet. These layers extract patterns from the 2D signal, and the resulting embedding is then passed through three fully connected layers to finally output the predicted relative joint angles of the human body, corresponding to the person's pose at time step t . The relative joint angles are the relative angle transformation of the distal joint with respect to the proximal joint represented in the axis angle format. The network was trained using stochastic gradient descent to minimize the mean squared error between the predicted joint angles and the ground truth.

We trained the network with a batch size of 128 and with a learning rate of 0.01 using the Adam optimizer. This was a design choice to predict the pose in the joint angle space instead of the position space, but the code provided can be easily modified to predict the pose in position space. The collected dataset also contains positions of the different joints that could be used to train such a model.

Evaluation. The model is validated with the test dataset, including various exercises. The ability of the model to infer each of the joint angles is illustrated in Supplementary Fig. 13. The model can infer the motion of the lower body better than the motion of the upper body. It is also found that the motion of the hands does not induce a systematic change in the tactile pressure map, and hence their joint angles are more difficult to infer. There was no additional loss used to train the system to be temporally consistent. The network figured out the correspondence between the trajectory and the pose, and smoothly transitioned from one pose to another accordingly. More results are provided in Supplementary Fig. 12 and Supplementary Video 5. We also evaluated how the performance of the model changes with a systematic reduction in the resolution on each of the feet from 32×32 to 1×1 (with the same ablation method used for letter classification).

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

Data that support the findings of this study are available from the corresponding authors upon reasonable request. Source data are provided with this paper.

Code availability

The code used to generate the plots within this paper is available from the corresponding authors upon reasonable request.

Received: 26 August 2020; Accepted: 18 February 2021;
Published online: 24 March 2021

References

- Dahiya, R. S., Metta, G., Valle, M. & Sandini, G. Tactile sensing—from humans to humanoid. *IEEE Trans. Robot.* **26**, 1–20 (2009).
- Winter, D. A. Human balance and posture control during standing and walking. *Gait Posture* **3**, 193–214 (1995).
- Someya, T., Bao, Z. & Malliaras, G. G. The rise of plastic bioelectronics. *Nature* **540**, 379–385 (2016).
- Yousef, H., Boukallel, M. & Althoefer, K. Tactile sensing for dexterous in-hand manipulation in robotics—a review. *Sens. Actuator A Phys.* **167**, 171–187 (2011).
- Yang, G. Z. et al. The grand challenges of Science Robotics. *Sci. Robot* **3**, eaar7650 (2018).
- Sundaram, S., Kellnhofer, P., Zhu, J. Y., Torralba, A. & Matusik, W. Learning the signatures of the human grasp using a scalable tactile glove. *Nature* **569**, 698–702 (2019).
- Poupyrev, I. et al. Project Jacquard: interactive digital textiles at scale. In *Proc. 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)* 4216–4227 (ACM, 2016).
- Zeng, W. et al. Fiber-based wearable electronics: a review of materials, fabrication, devices and applications. *Adv. Mater.* **26**, 5310–5336 (2014).
- Yan, W. et al. Advanced multimaterial electronic and optoelectronic fibers and textiles. *Adv. Mater.* **31**, 1802348 (2019).
- Rogers, J. A., Someya, T. & Huang, Y. Materials and mechanics for stretchable electronics. *Science* **327**, 1603–1607 (2010).
- Engler, A. J., Sen, S., Sweeney, H. L. & Discher, D. E. Matrix elasticity directs stem cell lineage specification. *Cell* **126**, 677–689 (2006).
- Moeslund, T. B., Hilton, A. & Kruger, V. A survey of advances in vision-based human motion capture and analysis. *Comput. Vis. Image Underst.* **104**, 90–126 (2006).
- Rein, M. et al. Diode fibres for fabric-based optical communications. *Nature* **560**, 214–218 (2018).
- Kim, D. H. et al. Stretchable and foldable silicon integrated circuits. *Science* **320**, 507–511 (2008).
- Wang, S., Oh, J. Y., Xu, J., Tran, H. & Bao, Z. Skin-inspired electronics: an emerging paradigm. *Acc. Chem. Res.* **51**, 1033–1045 (2018).
- Xenoma; <https://xenoma.com/#products/>
- Ahn, B. Y. et al. Omnidirectional printing of flexible, stretchable and spanning silver microelectrodes. *Science* **323**, 1590–1593 (2009).
- Truby, R. L. & Lewis, J. A. Printing soft matter in three dimensions. *Nature* **540**, 371–378 (2016).
- Adams, J. A. A closed-loop theory of motor learning. *J. Mot. Behav.* **3**, 111–150 (1971).
- Wang, Z., Yang, Z. & Dong, T. A review of wearable technologies for elderly care that can accurately track indoor position, recognize physical activities and monitor vital signs in real time. *Sensors* **17**, 341 (2017).
- Narayanan, V., Wu, K., Yuksel, C. & McCann, J. Visual knitting machine programming. *ACM Trans. Graph.* **38**, 63 (2019).
- Shima Seiki. Sds-one apex3 http://www.shimaseiki.com/product/design/sdsone_apex/flat/ (2011).
- Briseno, A. L. et al. Patterning organic single-crystal transistor arrays. *Nature* **444**, 913–917 (2006).
- Khang, D. Y., Jiang, H., Huang, Y. & Rogers, J. A. A stretchable form of single-crystal silicon for high-performance electronics on rubber substrates. *Science* **311**, 208–212 (2006).
- Ulyanov, D., Vedaldi, A. & Lempitsky, V. Deep image prior. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR '2018)* 9446–9454 (IEEE, 2018).
- D'Alessio, T. Measurement errors in the scanning of piezoresistive sensors arrays. *Sens. Actuator A Phys.* **72**, 71–76 (1999).
- Maaten, L. V. D. & Hinton, G. Visualizing data using *t*-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
- Lederman, S. J. & Klatzky, R. L. Haptic perception: a tutorial. *Atten. Percept. Psychophys.* **71**, 1439–1459 (2009).
- Bauby, C. E. & Kuo, A. D. Active control of lateral balance in human walking. *J. Biomech.* **33**, 1433–1440 (2000).
- Scott, J. et al. From kinematics to dynamics: estimating center of pressure and base of support from video frames of human motion. Preprint at <https://arxiv.org/abs/2001.00657> (2020).
- Okamura, A. M., Smaby, N. & Cutkosky, M. R. An overview of dexterous manipulation. In *Proc. 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings Vol. 1*, 255–262 (IEEE, 2000).
- Cheng, G. et al. A comprehensive realization of robot skin: sensors, sensing, control and applications. *Proc. IEEE* **107**, 2034–2051 (2019).
- Cho, K. et al. Learning phrase representations using RNN encoder–decoder for statistical machine translation. Preprint at <https://arxiv.org/abs/1406.1078> (2014).

Acknowledgements

This work is supported by the Toyota Research Institute. We thank L. Makatura, P. Kellnhofer, A. Kaspar and S. Sundaram for their helpful suggestions for this work. We also thank D. Rus for the use of the mechanical tester and J. L. McCann for providing us with the necessary code to programmatically work with our industrial knitting machine and visualize the knitting structure.

Author contributions

Y. Luo, W.S. and M.F. developed and implemented the functional fibre fabrication set-up. Y. Luo and W.S. conceived and implemented the sensor design, and performed

characterizations. Y. Luo and K.W. designed and fabricated full-body sensing textiles. Y. Luo, W.S., Y. Li, P.S., K.W. and B.L. conducted data collection. Y. Li conducted the self-supervised sensing correction. P.S. conducted the experiments on 3D pose prediction from the tactile footprints. Y. Li and P.S. implemented the classification framework. T.P., A.T. and W.M. supervised the work. All authors contributed to the study concept, conceived of the experimental methods, discussed and generated the results, and prepared the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41928-021-00558-0>.

Correspondence and requests for materials should be addressed to Y.Li, W.S., A.T. or W.M.

Peer review information *Nature Electronics* thanks Sihong Wang, Jun Chen and Lining Yao for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2021, corrected publication 2021

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- ☐ ☒ The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- ☐ ☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- ☒ ☐ The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- ☒ ☐ A description of all covariates tested
- ☒ ☐ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- ☐ ☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- ☒ ☐ For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- ☒ ☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- ☒ ☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- ☒ ☐ Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection Tactile frames from on-body human recording were captured using custom developed electronic module and commands. The poses and motions of human were captured with a commercial motion capture system, Xsens MVN.

Data analysis All data visualization and analysis were performed with customized codes based on open source libraries, including PyTorch and OpenCv.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The data that support the plots within the paper and other findings of this study are available from the corresponding authors upon reasonable request.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	One consenting female subject with multiple recordings
Data exclusions	No data excluded
Replication	Different devices made with the same materials were tested for multiple times with different activities.
Randomization	Subject was selected within the research group
Blinding	Data from the tests were analyzed by different authors

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☐	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☐	☒ Human research participants
☒	☐ Clinical data
☒	☐ Dual use research of concern

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Eukaryotic cell lines

Policy information about [cell lines](#)

Cell line source(s)	NA
Authentication	NA
Mycoplasma contamination	NA
Commonly misidentified lines (See ICLAC register)	NA

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics	Healthy female subject in the age range of 20-40
Recruitment	Consenting subjects were recruited from within the research group
Ethics oversight	NA

Note that full information on the approval of the study protocol must also be provided in the manuscript.