Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/143716/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Zhu, Zhongguan, Guo, Shihui, Qin, Yipeng ORCID: <https://orcid.org/0000-0002-1551-9126>, Chen, Xiaowei, Wu, Ronghui, Shi, Yating, Liu, Xiangyang and Liao, Minghong 2021. Robust elbow angle prediction with aging soft sensors via output-level domain adaptation. IEEE Sensors Journal 21 (20) , pp. 22976-22984. 10.1109/JSEN.2021.3091004 file

Publishers page: <https://doi.org/10.1109/JSEN.2021.3091004>
<<https://doi.org/10.1109/JSEN.2021.3091004>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.

Robust Elbow Angle Prediction with Aging Soft Sensors via Output-Level Domain Adaptation

Zhongguan Zhu, Shihui Guo, Yipeng Qin , Xiaowei Chen, Ronghui Wu, Yating Shi, Xiangyang Liu, Minghong Liao

Abstract—Wearable devices equipped with soft sensors provide a promising solution for body movement monitoring. Specifically, body movements like elbow flexion can be captured by monitoring the stretched soft sensors' resistance changes. However, in addition to stretching, the resistance of a soft sensor is also influenced by its aging, which makes the resistance a less stable indicator of the elbow angle. In this paper, we leverage the recent progress in Deep Learning and address the aforementioned issue by formulating the aging-invariant prediction of elbow angles as a *domain adaption* problem. Specifically, we define the soft sensor data (i.e., resistance values) collected at different aging levels as different domains and adapt a regression neural network among them to learn domain-invariant features. However, unlike the popular pairwise domain adaptation problem that only involves one source and one target domain, ours is more challenging as it has “infinite” target domains due to the non-stop aging. To address this challenge, we novelly propose an output-level domain adaptation approach which builds on the fact that the elbow angles are in a fixed range regardless of aging. Experimental results show that our method enables robust and accurate prediction of elbow angles with aging soft sensors, which significantly outperforms supervised learning ones that fail to generalize to aged sensor data.

Index Terms—elbow angle prediction, soft sensor, sensor aging, deep learning, transfer learning, domain adaptation

I. INTRODUCTION

WEARABLE devices targeted at body movement monitoring (BMM) have been applied in various fields of human-centered computing including human-computer interaction (HCI) [1] [2], healthcare [3] [4], and even military [5]. In these applications, wearable devices are usually designed with gloves or joint pads to capture the bending of body parts like fingers, arms, or legs. To facilitate long-term use, soft sensors have become the preferred choice for such wearable devices as they have almost no influence on the feels of clothing fabrics and thus improve the clothing comfort. In these devices, soft sensors are coiled on the cloth fibers and stretch with body movements. Such stretches change the resistance of soft sensors, commonly used as the *indicator* of body movements. With paired data of soft sensor signals and body movements collected, it is straightforward to estimate the mapping in-between by supervised learning [1].

This work was supported in part by the Fundamental Research Funds for the Central Universities (20720190006), the National Natural Science Foundation of China (61702433, 62072383) and the Cardiff/Xiamen Mobility Funds. The associate editor coordinating the review of this article and approving it for publication was Prof. Chao Tan. (Corresponding author: Yipeng Qin.)

Zhongguan Zhu (email: dosia96@163.com), Shihui Guo (email: guoshihui@xmu.edu.cn), Xiaowei Chen, Ronghui Wu, Yating Shi, Xiangyang Liu, Minghong Liao are with Xiamen University, Xiamen 361005, China.

Yipeng Qin (email: qiny16@cardiff.ac.uk) is with School of Computer Science and Informatics, Cardiff University, Cardiff CF10 3AT, UK.

Although promising, the aforementioned wearable devices suffer from the aging of soft sensors, making the learnt mapping invalid after certain amounts of usage. To delay or mitigate the aging effects, two techniques are employed by researchers and practitioners as common practice: replacement and using anti-aging materials. Between them, replacement is naive but widely used. Although effective in some cases, it does not apply to soft sensors as they are coiled around cloth fibres and thus cannot be replaced without damaging the wearable device. Unlike replacement, anti-aging materials have already been used as the coating for our soft sensors. While from our tests, the aging of the resulting soft sensors is still significant during a relatively short period (approx. 2 days). Thus, additional efforts are required to mitigate the aging effects of soft sensors.

In this work, we take elbow angle prediction as an example and address the aging problem by incorporating *domain adaptation* into the learning. Specifically, instead of naively assuming the prediction to be carried out in a single domain (i.e. supervised learning), we observed the drift of soft sensor signals along with its aging and defined the data collected at different aging levels of soft sensors as different domains. In this way, we convert the prediction of elbow angles with aging soft sensors to a domain adaptation problem that aims to learn a domain/aging-invariant regression neural network. Note that this formulation is not trivial as the non-stop aging of soft sensors results in “infinite” domains for adaptation, which is

different from the cases in most literature that only involves one source and one target domain [6] [7] [8]. Addressing this challenge, we utilize the fact that the elbow angles are in a fixed range regardless of the aging and propose an *output-level* domain adaptation method that involves two strategies: i) To make the elbow angles identically distributed, we intentionally asked the volunteers to do full elbow flexion at uniform speeds; ii) To facilitate statistical discrepancy estimation, we conduct domain adaptation at the low-dimensional output-level (*i.e.* a scalar) of the predictor instead of its high dimensional feature-level. That is, we apply domain adaptation to the output *angle scalars* of the predictor instead of the output *feature vectors* of its intermediate layers. This is critical as the low dimensionality of the predictor output guarantees the validity of statistical measures which suffer from the notorious *curse of dimensionality* [9]. Experimental results show that our method enables robust and accurate prediction of elbow angles with aging soft sensors, which significantly outperforms supervised learning ones that fail to generalize to aged sensor data.

Contributions. First, we pioneer the aging-invariant prediction of elbow angles using domain adaptation, which reduces the mean prediction errors by half from 16.37° to 8.78° . Second, we propose an output-level domain adaption approach and two associated strategies, which outperforms feature-level ones in both accuracy and robustness.

II. RELATED WORK

Body Movement Monitoring (BMM). From its fundamental nature, BMM works as an important building block in various applications [10]–[17]. To date, the dominating approaches of BMM fall into two categories: vision-based ones and those using wearable devices. Between them, the vision-based methods [18]–[20] detect human body movements by analyzing the videos captured by RGB or infrared cameras. Although being cost-effective, these methods are sensitive to optical conditions (*e.g.* occlusion, lighting) and thus struggled to be applied in complex real-world scenarios. To make BMM more robust, wearable devices were employed. For example, Huang *et al.* [21] proposed a real-time human pose reconstruction method by training a deep neural network with the signals captured from 6 Inertial Measurement Units (IMUs). Eyobu *et al.* [22] addressed the data scarcity problem of IMU-based BMM by incorporating novel feature representation and data augmentation methods. Despite their effectiveness, IMUs are “foreign bodies” to clothes and thus inevitably reduce the clothing comfort. To this end, people turned to soft sensor-based wearable devices that are indistinguishable from daily clothing. For example, Farrington *et al.* [23] pioneered to use of an advanced soft sensor to measure upper limb and body movements. Wang *et al.* [24] successfully captured finger, elbow, and other body movements using corresponding wearable devices equipped with their composite yarn. Glauser *et al.* [1] designed a stretch-sensing soft glove and used it to interactively estimate hand poses with the aid of a deep neural network. We refer interested readers to [25] for a detailed review. In this work, we follow this trend and investigate the

BMM using soft sensor-based wearable devices.

Soft Sensors. As aforementioned, soft sensors have become a popular choice for the wearable devices targeted at BMM from their high biocompatibility, ductility, and portability [26]–[30]. Although promising, soft sensors suffer from the aging problem, which impedes their application in consumer products. In addition to the common practice aimed at mitigating sensor aging (*e.g.* anti-oxidation coating), Wang *et al.* [31] designed a new sensor that can measure the degree of aging for crystal vibrating materials and proposed a method to compensate the effects of aging accordingly. However, their method relies on the measurement provided by their new sensor and is thus only applicable to crystal vibrating materials. By investigating the origins of soft sensor aging, Scilingo *et al.* [32] proposed a method to compensate for the aging effects caused by plastic deformation. However, their method is dedicated to plastic deformation and is thus not applicable to the aging caused by other factors. In this work, we address the abovementioned issues with the idea of *domain adaptation*, which is borrowed from the machine learning community. As a result, our method is independent of sensor materials, and implicitly deals with all aging factors simultaneously.

Domain Adaptation. Being a sub-field of transfer learning, domain adaptation aims to mitigate the differences among domains, which enables a model trained in one or more source domains to be successfully applied to different target domains. Inspired by the seminal work of Pan *et al.* [33] and the recent progress in deep learning, massive amounts of effort have been made in domain adaptation. Yosinski *et al.* [34] first demonstrated the possibility of adapting a pre-trained AlexNet [35] to different domains by fine-tuning several of its last layers. However, fine-tuning requires additional labeled data in the target domain, which is not available for many applications. Addressing this issue, Tzeng *et al.* [36] proposed to add an adaptation layer and minimize the MMD loss between the feature distributions of the source and target domains at the adaptation layer. Their method is further extended by Long *et al.* [6] by incorporating a multi-kernel MMD loss and perform adaptation at multiple layers. For soft sensors, Liu *et al.* [37] applied domain adaption to improve soft sensor models used to predict product quality across the different modes of multi-grade chemical processes. To the best of our knowledge, we are the first to show that domain adaptation can also be used to mitigate the aging effects of soft sensors for body movement monitoring. Furthermore, we novelly utilize the low dimensionality of predicted elbow angles and propose to conduct domain adaption at the output level of the predictor. To our knowledge, this is rarely discussed in previous work as the outputs of many machine learning tasks (*e.g.* image classification) are high dimensional vectors on which there is no obvious benefit to applying output-level domain adaptation.

III. HARDWARE SPECIFICATIONS

This section shows the hardware specifications for our work. Specifically, we first introduce the details of the soft sensors used (Section III-A) and then discuss how they are sewed in the SmartPad for elbow angle prediction (Section III-B).

Finally, we show the hardware used to collect, transmit and process the soft sensor data (Section III-C).

A. Soft Sensors

To maximize the clothing comfort for long-term use, we employed silk fibers based soft sensors coiled on supporting yarns, which have a similar touch to common clothing materials. More specifically, the silk fibers are coated with conductive slurry to conduct electricity. The conductive slurry is made by mixing a polyurethane solution diluted to 6.5mg/mL and a silver nanowires solution diluted to 5mg/mL at the same volume ratio. The yarn tension sensor is composed of natural fiber or human-made fiber, conductive layer slurry and electrode material. The strain coefficient can reach 33.19% with high sensitivity and flexibility. The knitting structure is selected as the closed chain, the yarn was braided, and the silver nanowires with a concentration of 5.2mg/mL were dipped and dried. The braided silk braided chain structure was inverted on the surface. The sensor is manually sewn into the elbow pad through the flat stitching. The sensor is connected to the circuit by a sewn silver thread (a cotton fiber impregnated with silver ions).

B. SmartPad for Elbow Angle Prediction

Fig. 1: The SmartPad used in our work.

Our SmartPad for elbow angle prediction consists of six soft sensors evenly distributed around the elbow pad (see Fig. 5 SmartPad). It uses common clothing material as fabric material (Fig. 1), where the flexible sensors are sewn without influencing clothing comfort and blend perfectly with the elbow pads. When being used, sensor signals consisting of six scalar values will be collected and sent by the PCB equipped with Bluetooth.

C. Data Collection, Transmission and Processing

Fig. 2: Left: SmartPad and the reflective markers used in Qualisys Motion-Capture. Right: markers (green) and vectors (blue) used in ground truth elbow angle calculation.

Data Collection. Along with the flexible sensor signals that are captured by our SmartPad, we collect ground truth elbow angles by augmenting the SmartPad with reflective markers and reconstruct human body movement with motion capture techniques. As Fig. 2 shows, we determine the ground truth elbow angle values using the 3D positions of the markers (green) which are captured by our high-precision ($\epsilon = 0.07\text{mm}$) Miquis M5 - Qualisys motion-capture system¹. Specifically, we define the ground truth elbow angle as the angle θ between vectors $\overrightarrow{P_y P_x}$ and $\overrightarrow{Q_x Q_y}$, where $P_x = (P_1 + P_2)/2$, $P_y = (P_3 + P_4)/2$, $Q_x = (Q_1 + Q_2)/2$, $Q_y = (Q_3 + Q_4)/2$ are the midpoints of the corresponding pairs of markers respectively. Then, we calculate the ground truth elbow angle as:

$$\theta = \arccos \frac{\overrightarrow{P_y P_x} \cdot \overrightarrow{Q_x Q_y}}{\|\overrightarrow{P_y P_x}\| \cdot \|\overrightarrow{Q_x Q_y}\|} \quad (1)$$

To estimate the error incurred by the precision of the motion capture system, we first estimate the angle error for each of the two vectors $\overrightarrow{P_y P_x}$ and $\overrightarrow{Q_x Q_y}$. As Fig. 3 shows, we take $\overrightarrow{P_y P_x}$

Fig. 3: Illustration of error estimation.

as an example and set the x-axis of the coordinate system along $\overrightarrow{P_y P_x}$ with its origin at the midpoint of $P_x P_y$. Let P'_x and P'_y be the marker positions captured by our Qualisys system, we calculate the maximum error angle $\Delta\theta_P$ of $\overrightarrow{P_y P_x}$ as:

$$\Delta\theta_P = \max_{P'_x, P'_y} \arccos \frac{\overrightarrow{P_y P_x} \cdot \overrightarrow{P'_y P'_x}}{\|\overrightarrow{P_y P_x}\| \cdot \|\overrightarrow{P'_y P'_x}\|} \quad (2)$$

$$s.t. \|P_x - P'_x\| \leq \epsilon, \|P_y - P'_y\| \leq \epsilon$$

Geometrically, it can be observed that $\Delta\theta_P$ achieves its maximum when $P'_y P'_x$ goes through the origin and is tangent to both error spheres (see Fig. 3). Thus, we have

$$\Delta\theta_P = \arccos \frac{\sqrt{(L/2)^2 - \epsilon^2}}{L/2} \quad (3)$$

For $\overrightarrow{P_y P_x}$, we substitute $\epsilon = 0.07\text{mm}$ and $L = 350\text{mm}$ into Eq. 3 and have $\Delta\theta_P = 0.023^\circ$. Similarly, for $\overrightarrow{Q_x Q_y}$, we substitute $\epsilon = 0.07\text{mm}$ and $L = 200\text{mm}$ into Eq. 3 and have $\Delta\theta_Q = 0.040^\circ$. Assume both $\Delta\theta_P$ and $\Delta\theta_Q$ fully contribute to the final error, we approximately bound the error of ground truth elbow angle by $\Delta\theta \leq \Delta\theta_P + \Delta\theta_Q = 0.063^\circ$, which is negligible compared to that of our method ($> 8^\circ$).

Data Transmission. As aforementioned, the soft sensor signals captured by our SmartPad are transmitted by a PCB board

¹<https://www.qualisys.com/cameras/miquis/#tech-specs>, accessed 22/04/2021.

of six channels (Fig. 4), where the amplified voltage measurements are processed by a low-pass filter with a bandwidth of 300 Hz. The channels are selected by a multiplex voltage divider. We used a Wheatstone bridge structure calculation to distinguish the voltage and reference voltage $V_{ref} = V_{cc}/2$ of each sensor, where V_{cc} represents the voltage of the access circuit. The input voltage to the analog-digital conversion $V_{adc_{in}}$ is defined as (ignoring the effect of the low-pass filter):

$$V_{adc_{in}} = \left(\frac{V_{CC} * R_{sensor_i}}{R_i + R_{sensor_i}} - V_{CC}/2 \right) * Gain, \quad (4)$$

where R_i denotes the divider resistor, R_{sensor_i} indicates the resistance of the i^{th} soft sensor, and $Gain$ denotes the magnification factor of the amplifier unit.

Through the mold, electrical signals are converted into digital signals and transmitted via a bluetooth chip (Nrf51822 at 20 frames per second) to a mobile device (Huawei Mate 20, memory: 6GB, Android version: 9). The data are then transmitted and decoded at the server for later use.

Fig. 4: Circuit diagram of the PCB board in our SmartPad [38].

Data Processing. We processed the collected data with a server machine equipped with an 8-core I7 9700k CPU and an 8GB GeForce GTX 2070 GPU.

IV. LEARNING AN AGING-INVARIANT PREDICTOR

Let $\mathbf{x} = [R_a, R_b, R_c, R_d, R_e, R_f] \in \mathcal{X}$ be the signal vector collected by our SmartPad, $y \in \mathcal{Y}$ be its corresponding elbow angle, we aim to train a predictor which estimates the function $f: \mathcal{X} \rightarrow \mathcal{Y}$ by fitting a training dataset $\mathcal{D}_{train}$. However, due to the aging of soft sensors, the data points in $\mathcal{D}_{train}$ are usually far from those in practice. As a result, although being relatively easy to achieve a high prediction accuracy on $\mathcal{D}_{train}$, the predictor trained with supervised learning is prone to overfitting and worsens significantly on the test data collected from the continuously aging soft sensors in real-world scenarios. To this end, we borrow the ideas from domain adaptation [6] [7] [8] and propose to explicitly regularize the predictor to learn aging-invariant features during training. Specifically, we define the soft sensor data $\mathbf{x}$ collected at different aging levels as different domains and adapt the predictor among them to learn domain-invariant features. However, this is not trivial because the pairwise domain adaptation problems studied by most existing works [6] [7] [8] only require the

predictor to be adapted to a single target domain. In contrast, ours requires the predictor to be adapted to “infinite” target domains cause by the non-stop aging of soft sensors.

Addressing this challenge, we propose two strategies based on the invariance in the range of elbow angles (*i.e.* $[30^\circ, 180^\circ]$), which is fixed regardless of aging: 1) During the collection of $\mathcal{D}_{train}$, we intentionally asked the volunteers to do full elbow flexion at uniform speeds so that the elbow angles are guaranteed to be identically distributed; 2) Facilitating the statistical discrepancy estimation, we conduct domain adaptation on the low-dimensional output scalars of the predictor instead of the high-dimensional (*e.g.* 512) output feature vectors of intermediate layers. Fig. 5 shows an overview of our method.

A. Curated Data Collection

To make sure that the elbow angles in $\mathcal{D}_{train}$ are identically distributed throughout the data collection process, we first recruited 10 volunteers of different characteristics (*e.g.* gender, height, weight, etc.) and asked them to wear our SmartPad casually; Then, we ask the volunteers to do full elbow flexion at uniform speeds (Fig. 6). That is, the volunteers are asked to repeatedly bend their elbows to approximately 30° and extend them to approximately 180° . Note that we do not require all volunteers to do elbow flexion at the same speed but ask them to do them evenly (*i.e.* uniform speeds). The collected elbow angle data are justified to be approximately identically distributed by visual inspection on their distributions.

In this way, we obtained our training dataset $\mathcal{D}_{train} = (\mathbf{x}_i, y_i)_{i=1}^n$, where n is the number of data points. Note that the soft sensors are continuously aging during data collection.

B. Output-Level Domain Adaptation

Network Architecture. Leveraging the power of deep learning, we build the predictor with a neural network of 9 fully-connected layers wrapped with Leaky ReLU ($\alpha = 0.01$) [39] activation functions, respectively. Details of the network architecture are shown in Fig. 5.

Training. As above-mentioned, we collected a training dataset $\mathcal{D}_{train} = (\mathbf{x}_i, y_i)_{i=1}^n$ for the training of our elbow angle predictor. However, these data are collected from the continuously aging soft sensors, making the training prone to overfitting. Addressing this problem, we split $\mathcal{D}_{train}$ into two halves $\mathcal{D}_a = (\mathbf{x}_i, y_i)_{i=1}^{n/2}$ and $\mathcal{D}_b = (\mathbf{x}_i, y_i)_{i=n/2+1}^n$ of different degrees of aging, and encourage the learning of domain-invariant features by performing *output-level* domain adaption between them. Specifically, we aim to fulfill two objectives:

- Minimizing the prediction error on $\mathcal{D}_{train}$. Let $\hat{y}$ be the predicted elbow angle, and we measure the prediction error by a mean square error (MSE) loss:

$$L_{mse} = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 \quad (5)$$

- Minimizing the statistical distance between the two output distributions of the predictor on $\mathcal{D}_a$ and $\mathcal{D}_b$. Following common practice [6] [7] [8], we use maximum mean

Fig. 5: Overview of the proposed method. Our neural network based predictor takes soft sensor signals $[R_a, R_b, R_c, R_d, R_e, R_f]$ captured by the SmartPad as input and predicts elbow angles $\hat{y}$ accordingly. The model is trained by minimizing a combination of MSE and MMD loss terms weighted by λ at output-level to approximate ground truth elbow angles y .

Fig. 6: Curated data collection. The volunteers were asked to do full elbow flexion at uniform speeds.

discrepancy (MMD) to measure the distance. As discussed in [40], observing samples from two distributions P and Q , MMD can be used as a kernel two-sample test of the null hypothesis $P = Q$. In a nutshell, MMD represents the distances between two distributions as the distances between the *mean embeddings* of them in a reproducing kernel Hilbert space. Its key idea is that the *mean embeddings* of P and Q are the same if and only if $P = Q$. We refer interested audiences to [40] for a rigorous and detailed proof. In practice, our MMD loss can be empirically estimated or calculated as [40]:

$$\begin{aligned}
 L_{\text{mmd}} &= \|\mathbb{E}[\phi(\hat{Y}_a)] - \mathbb{E}[\phi(\hat{Y}_b)]\|_{\mathcal{H}}^2 \\
 &= \frac{1}{m^2} \sum_{i,j=1}^m k(\hat{y}_i^a, \hat{y}_j^a) - \frac{2}{mn} \sum_{i,j=1}^{m,n} k(\hat{y}_i^a, \hat{y}_j^b) \\
 &\quad + \frac{1}{n^2} \sum_{i,j=1}^n k(\hat{y}_i^b, \hat{y}_j^b)
 \end{aligned} \quad (6)$$

where $\hat{Y}_a$ and $\hat{Y}_b$ denote the outputs of the predictor on $\mathcal{D}_a$ and $\mathcal{D}_b$ respectively, $\hat{y}_i^a \in \hat{Y}_a$ ($i = 1, 2, \dots, m$) and $\hat{y}_i^b \in \hat{Y}_b$ ($i = 1, 2, \dots, n$) are samples of $\hat{Y}_a$ and $\hat{Y}_b$, $\phi(Y)$ is the feature embedding of Y in a reproducing kernel Hilbert space $\mathcal{H}$. Moreover, we follow [41] and use a multi-kernel MMD which averages the MMD over 5 different Gaussian kernels $k(y^a, y^b) = \exp(-\|y^a - y^b\|^2 / \gamma_t)$, where $t = \{-2, -1, 0, 1, 2\}$. Specifically, $\gamma_t = 2^t \gamma$ where we follow the common *median heuristic* [41] and set the bandwidth γ to be the median distance between points in the aggregate sample. Note that the validity of MMD minimization is ensured by i) our curated data collection process (Section IV-A), which guarantees that the outputs of the predictor are identically distributed; and ii) the low dimensionality of $\hat{Y}_a$ and $\hat{Y}_b$, which prevents the problematic drop (at best polynomial) of MMD values with dimension that makes them invalid [9].

Between them, L_{mse} is a standard loss term used in supervised learning and L_{mmd} is the loss term for our output-level domain adaptation, which encourages the learning of an aging-invariant predictor by regularizing its output distributions to be the same regardless of the aging. Thus, the overall loss function of our method is:

$$L = L_{\text{mse}} + \eta \cdot \lambda \cdot L_{\text{mmd}} \quad (7)$$

where λ is a weighting parameter, η is a learning rate decaying parameter that prioritizes the minimization of L_{mse} in the early stage of the training [42].

Inference. During inference, we predict the elbow angle $\hat{y}$ by simply feeding the soft sensor signal vector $\mathbf{x}$ into the predictor and making a forward pass. This is highly efficient because a forward pass only takes a few milliseconds. Thus, our predictor is not the bottleneck of the system as it has a frequency (20 kHz) much higher than the sampling frequency (20 Hz) of the SmartPad.

V. EXPERIMENTAL RESULTS

Empirically, we justify the effectiveness and rationale of our method by the following experimental results.

A. Experimental Setup

Datasets. We collected 60,000 pieces of data, *i.e.* $(\mathbf{x}_j, y_j)_{j=1}^m$ and $m = 60,000$, at different time slots spanning two days, which corresponds to different degrees of soft sensor aging. Then, we split the data into the training dataset $\mathcal{D}_{train}$, the validation dataset $\mathcal{D}_{val}$ and the test dataset $\mathcal{D}_{test}$ as follows:

- $\mathcal{D}_{train} = \mathcal{D}_a \cup \mathcal{D}_b$ where $\mathcal{D}_a = (\mathbf{x}_i, y_i)_{i=1}^{9,000}$ and $\mathcal{D}_b = (\mathbf{x}_i, y_i)_{i=10,001}^{19,000}$. That is, $\mathcal{D}_{train}$ contains 90% of the first 20,000 pieces of data, while the other 10% are included in $\mathcal{D}_{val}$ (see below).
- $\mathcal{D}_{val} = (\mathbf{x}_i, y_i)_{i=9,001}^{10,000} \cup (\mathbf{x}_i, y_i)_{i=19,001}^{20,000}$.
- $\mathcal{D}_{test} = (\mathbf{x}_i, y_i)_{i=20,001}^{60,000}$

Our split of $\mathcal{D}_{train}$ and $\mathcal{D}_{val}$ matches the idea of aging-invariant learning as it maximized the difference of soft sensor aging between the two datasets, which enables better early stopping [43] and thus helps to avoid overfitting. Note that we use a large $\mathcal{D}_{test}$ consisting of data points from different levels of aging to demonstrate the superiority of our method in accurately predicting elbow angles against sensor aging.

Metrics. We evaluate our method by the average of the errors in the test dataset $\mathcal{D}_{test}$:

$$err = \frac{1}{40,000} \sum_{i=20,001}^{60,000} |\hat{y}_i - y_i| \quad (8)$$

where $\hat{y}_i$ denotes the predicted elbow angle and y_i denotes the ground truth elbow angle in $\mathcal{D}_{test}$. Note that we run all experiments eight times to get a better understanding of their mean and variance in the metrics.

Implementation Details. We trained the model with an SGD optimizer of learning rate $1e^{-5}$. We did not use any normalization layers or regularizers during training. We used a batch size of 100.

B. Domain Adaptation v.s. Supervised Learning

To show the effectiveness of domain adaptation, we compare our method's prediction accuracy against a baseline model on $\mathcal{D}_{test}$. To facilitate a fair comparison, the baseline model is trained using the same network architecture, hyperparameters and datasets (*i.e.* $\mathcal{D}_{train}$ and $\mathcal{D}_{val}$) as our method but in a supervised learning way. In other words, the *only* difference between the baseline model and our method is that the baseline model is trained with the loss function $L_{base} = L_{mse}$ instead of Eq. 7. As Fig. 8 shows, our method significantly outperforms the baseline model in both the accuracy and the robustness of elbow angle prediction in the scenario of soft sensor aging. In general, our predictor achieves a mean error of 8.78° on $\mathcal{D}_{test}$, which approximately reduces that of the baseline model, *i.e.* 16.37° , by half. More specifically, we observed that the performance gap between our method and the baseline model, *i.e.* $|\hat{y}_i - y_i| - |\hat{y}_i^{base} - y_i|$, grows with i . Since our data are collected as time series, these increasing gaps justify that

incorporating domain adaptation indeed helps the learning of aging-invariant features. The discussion of outliers in Fig. 8 is detailed as follows.

Outliers. After carefully checking the data, we found that the outliers stem from meaningless resistance values caused by sensor faults.

Interestingly, we observed that our method also handles such outliers better than the baseline model (Fig. 7). This further demonstrates the superiority of our method.

Fig. 7: Histograms of the error distributions. Our method has much less outliers than the baseline. Note that the y -axis is in logarithmic scale.

C. Output-level v.s. Feature-level Domain Adaptation

To justify that our output-level domain adaptation works better than traditional feature-level ones, we compared the prediction errors of our method and those of its feature-level variants (Fig. 9). These feature-level variants exhaustively perform domain adaptation, *i.e.* minimizing the MMD (Eq. 7), at all intermediate layers of the network. Specifically, we name the variants LX where X denotes the layer number at which domain adaptation is performed. For example, $L1$ refers to the variant where the domain adaptation is carried out at layer 1 of the network. As Fig. 9 show, our method has the smallest mean and median, the lowest variance among trials and the least outliers, which is the most stable and accurate prediction model.

D. Algorithm Performance v.s. Datasets Splits

As Fig. 10 shows, we test the performance of both our method and the baseline on different splits of the datasets. Specifically, we split the datasets following the same paradigm described in Section V-A. Thus, for all splits, the validation set contains 2,000 samples and the test set contains $(60,000 - 2,000 - \text{Train set size})$ samples. It can be observed that our method outperforms the baseline on all splits, which justifies both the superiority of the proposed method and that the comparison results is not affected by the split of datasets.

E. Verification of Sensor Aging and its Impacts

In this section, we justify the motivation and rationale of our work by verifying that the flexible sensors in our SmartPad indeed aged and how such aging impacted the prediction of elbow angles.

Fig. 8: Comparison of predicted elbow angles and errors between the baseline and our method.

(a) Boxplots of our method and its feature-level variants. Let Q_1 and Q_3 be the lower and upper quartiles of the eight trials respectively, we follow standard practice and compute $IQR = Q_3 - Q_1$. The outliers are identified as the data outside range $[Q_1 - 1.5IQR, Q_3 + 1.5IQR]$.

Model	Ave-err	Var	Var w/o outliers
Baseline	16.37	627.09	146.94
Layer1	9.57	134.80	43.71
Layer2	12.72	4339.99	53.18
Layer3	11.39	1583.97	54.43
Layer4	13.38	3590.57	68.14
Layer5	9.49	639.52	36.34
Layer6	8.82	158.38	37.94
Layer7	9.95	117.16	41.59
Layer8	12.13	2598.69	50.07
Our method	8.78	95.86	35.35

(b) Quantitative metrics of our method and its feature-level variants.

Fig. 9: Comparison of our method against its feature-level variants.

Fig. 10: Boxplots of our method (blue) and the baseline (red) regarding their performance against different splits of datasets.

As a straightforward evidence of sensor aging, we test how the resistance of sensors changes over time when the SmartPad is stationary. Specifically, we first produce a new SmartPad and leave it in place. Then, we record the resistance of its six sensors every two hours. As Fig. 11 shows, the resistance of all six sensors changed significantly in a non-linear manner, which will undoubtedly impact the elbow angle prediction if not dealt with properly.

Furthermore, we measure sensor aging with the distance between the distributions of resistance values that are collected from non-aged and aged sensors respectively.

Applying such a measure to the sensor data we collected, i.e. $\mathcal{R} = (\mathbf{x}_j)_{j=1}^m$ where $m = 60,000$, we first split it into

Fig. 11: Changes of sensor resistance over time for a stationary SmartPad.

Fig. 12: Verification of sensor aging. We use the MMD between the distribution of non-aged sensor data $\mathcal{R}_0$ and that of aged sensor data $\mathcal{R}_i (i \neq 0)$ as the indicator of sensor aging. The red box highlights the region where $\mathcal{R}_i$ is close to $\mathcal{R}_0$.

small portions of 1,000 samples each as,

$$\mathcal{R} = \mathcal{R}_0 \cup \mathcal{R}_1 \cup \mathcal{R}_2 \cup \dots \cup \mathcal{R}_{59} \quad (9)$$

where $\mathcal{R}_i = (\mathbf{x}_j)_{j=i*1000}^{(i+1)*1000-1}$. Then, in consistency with the proposed method, we compute the MMD between $\mathcal{R}_i (i \neq 0)$ and $\mathcal{R}_0$, and visualize the results in Fig. 12.

As Fig. 12 shows, it can be observed that the flexible sensors in our SmartPad indeed aged and that such aging significantly changed the distributions of sensor data against usage. In addition, we observed that the impacts of aging on sensor data distributions are complex. Despite their upward trend, the MMD values were oscillating and can sometimes be very low, i.e. the aged sensors have similar distributions to the non-aged one. On one hand, this indicates the complex distributional shifts of input data that are challenging for supervised learning, which justifies our use of domain adaptation. On the other hand, this explains the abnormal behaviour of the Baseline model in Fig. 8: the improvement of its performance in intervals $[25500, 30000]$ and $[33000, 36000]$ is not from the superiority of the model but the changes of sensor data distributions.

VI. CONCLUSION

In this paper, we proposed a novel domain adaptation method for aging-invariant elbow angle prediction. Instead of naively applying existing domain adaptation methods that are mostly carried out at feature-level, we proposed to perform domain adaptation at the output-level of the prediction network, which improves both its robustness and accuracy.

Limitations and Future Work. Although effective, there are still several limitations of our method. First, our method

cannot fully eliminate the outliers caused by sensor faults. Second, our output-level domain adaptation method is more suitable for applications with low-dimensional outputs, whose performance may drop with the ones with high-dimensional outputs. In future work, we will address these limitations and explore the interpretation of the learned model.

REFERENCES

- [1] O. Glauser, S. Wu, D. Panozzo, O. Hilliges, and O. Sorkine-Hornung, "Interactive hand pose estimation using a stretch-sensing soft glove," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, pp. 1–15, 2019.
- [2] Inhyuk Moon, Myungjoon Lee, Junuk Chu, and Museong Mun, "Wearable emg-based hci for electric-powered wheelchair users with motor disabilities," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005, pp. 2649–2654.
- [3] S. Brewster, J. Lumsden, M. Bell, M. Hall, and S. Tasker, "Multimodal eyes-free interaction techniques for wearable devices," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2003, pp. 473–480.
- [4] C. C. Y. Poon, Y. M. Wong, and Y. Zhang, "M-health: The development of cuff-less and wearable blood pressure meters for use in body sensor networks," in *2006 IEEE/NLM Life Science Systems and Applications Workshop*, 2006, pp. 1–2.
- [5] H. Lee, J. Tak, and J. Choi, "Wearable antenna integrated into military berets for indoor/outdoor positioning system," *IEEE Antennas and Wireless Propagation Letters*, vol. 16, pp. 1919–1922, 2017.
- [6] M. Long, Y. Cao, J. Wang, and M. Jordan, "Learning transferable features with deep adaptation networks," in *International conference on machine learning*. PMLR, 2015, pp. 97–105.
- [7] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance," *arXiv preprint arXiv:1412.3474*, 2014.
- [8] Y. Zhu, F. Zhuang, J. Wang, J. Chen, Z. Shi, W. Wu, and Q. He, "Multi-representation adaptation network for cross-domain image classification," *Neural Networks*, vol. 119, pp. 214–221, 2019.
- [9] A. Ramdas, S. J. Reddi, B. Póczos, A. Singh, and L. Wasserman, "On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [10] L. Chen, J. Hoey, C. D. Nugent, D. J. Cook, and Z. Yu, "Sensor-based activity recognition," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 42, no. 6, 2012.
- [11] D. Gavrilu, "The visual analysis of human movement: A survey," *Computer Vision and Image Understanding*, vol. 73, no. 1, 1999.
- [12] R. Poppe, "A survey on vision-based human action recognition," *Image and Vision Computing*, vol. 28, no. 6, pp. 976 – 990, 2010.
- [13] K. Aminian and B. Najjafi, "Body movement monitoring system and method," Nov. 28 2006, US Patent 7,141,026.
- [14] T. Choudhury, G. Borriello, S. Consolvo, D. Haehnel, B. Harrison, B. Hemingway, J. Hightower, P. P. Klasnja, K. Koscher, A. LaMarca, J. A. Landay, L. LeGrand, J. Lester, A. Rahimi, A. Rea, and D. Wyatt, "The mobile sensing platform: An embedded activity recognition system," *IEEE Pervasive Computing*, vol. 7, no. 2, pp. 32–41, 2008.
- [15] K. Van Laerhoven and K. Aidoo, "Teaching context to applications," *Personal and Ubiquitous Computing*, vol. 5, no. 1, pp. 46–49, 2001.
- [16] C. R. Wren and E. M. Tapia, "Toward scalable activity recognition for sensor networks," in *Location- and Context-Awareness*, 2006.
- [17] M. Philipose, K. P. Fishkin, M. Perkowitz, D. J. Patterson, D. Fox, H. Kautz, and D. Hahnel, "Inferring activities from interactions with objects," *IEEE Pervasive Computing*, vol. 3, no. 4, pp. 50–57, 2004.
- [18] Z. Cao, G. H. Martinez, T. Simon, S.-E. Wei, and Y. A. Sheikh, "Openpose: realtime multi-person 2d pose estimation using part affinity fields," *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [19] X. Chen and A. L. Yuille, "Articulated pose estimation by a graphical model with image dependent pairwise relations," in *Advances in neural information processing systems*, 2014, pp. 1736–1744.
- [20] C.-J. Chou, J.-T. Chien, and H.-T. Chen, "Self adversarial training for human pose estimation," in *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, 2018.
- [21] Y. Huang, M. Kaufmann, E. Aksan, M. J. Black, O. Hilliges, and G. Pons-Moll, "Deep inertial poser: learning to reconstruct human pose from sparse inertial measurements in real time," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 6, pp. 1–15, 2018.
- [22] O. Steven Eyobu and D. S. Han, "Feature representation and data augmentation for human activity classification based on wearable imu sensor data using a deep lstm neural network," *Sensors*, vol. 18, no. 9, 2018.
- [23] J. Farrington, A. J. Moore, N. Tilbury, J. Church, and P. D. Biemond, "Wearable sensor badge and sensor jacket for context awareness," in *Digest of Papers. Third International Symposium on Wearable Computers*. IEEE, 1999, pp. 107–113.
- [24] Z. Wang, Y. Huang, J. Sun, Y. Huang, H. Hu, R. Jiang, W. Gai, G. Li, and C. Zhi, "Polyurethane/cotton/carbon nanotubes core-spun yarn as high reliability stretchable strain sensor for human motion detection," *ACS applied materials & interfaces*, vol. 8, no. 37, 2016.
- [25] M. Amjadi, K.-U. Kyung, I. Park, and M. Sitti, "Stretchable, skin-mountable, and wearable strain sensors and their potential applications: a review," *Advanced Functional Materials*, vol. 26, no. 11, 2016.
- [26] W. Li, F. Xu, W. Liu, Y. Gao, K. Zhang, X. Zhang, and Y. Qiu, "Flexible strain sensor based on aerogel-spun carbon nanotube yarn with a core-sheath structure," *Composites Part A: Applied Science and Manufacturing*, vol. 108, pp. 107–113, 2018.
- [27] G. Chen, H. Wang, R. Guo, M. Duan, Y. Zhang, and J. Liu, "Superelastic egain composite fibers sustaining 500% tensile strain with superior electrical conductivity for wearable electronics," *ACS Applied Materials & Interfaces*, vol. 12, no. 5, pp. 6112–6118, 2020.
- [28] L. Wang, L. Wang, Y. Zhang, J. Pan, S. Li, X. Sun, B. Zhang, and H. Peng, "Weaving sensing fibers into electrochemical fabric for real-time health monitoring," *Advanced Functional Materials*, vol. 28, no. 42, p. 1804456, 2018.
- [29] Z. He, G. Zhou, J.-H. Byun, S.-K. Lee, M.-K. Um, B. Park, T. Kim, S. B. Lee, and T.-W. Chou, "Highly stretchable multi-walled carbon nanotube/thermoplastic polyurethane composite fibers for ultrasensitive, wearable strain sensors," *Nanoscale*, vol. 11, no. 13, 2019.
- [30] Y. Cheng, R. Wang, J. Sun, and L. Gao, "A stretchable and highly sensitive graphene-based fiber for sensing tensile strain, bending, and torsion," *Advanced materials*, vol. 27, no. 45, pp. 7365–7371, 2015.
- [31] X. Wang, L. Winemberg, D. Su, D. Tran, S. George, N. Ahmed, S. Palosh, A. Dobin, and M. Tehranipoor, "Aging adaption in integrated circuits using a novel built-in sensor," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 1, 2014.
- [32] Enzo Pasquale Scilingo, F. Lorussi, A. Mazzoldi, and D. De Rossi, "Strain-sensing fabrics for wearable kinaesthetic-like systems," *IEEE Sensors Journal*, vol. 3, no. 4, pp. 460–467, 2003.
- [33] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang, "Domain adaptation via transfer component analysis," *IEEE Transactions on Neural Networks*, vol. 22, no. 2, pp. 199–210, 2010.
- [34] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Advances in neural information processing systems*, 2014, pp. 3320–3328.
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, vol. 25, 2012.
- [36] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell, "Deep domain confusion: Maximizing for domain invariance," *arXiv preprint arXiv:1412.3474*, 2014.
- [37] Y. Liu, C. Yang, K. Liu, B. Chen, and Y. Yao, "Domain adaptation transfer learning soft sensor for product quality prediction," *Chemometrics and Intelligent Laboratory Systems*, vol. 192, p. 103813, 2019.
- [38] H. Zhang, Z. Chen, S. Guo, J. Lin, Y. Shi, X. Liu, and Y. Ma, "Sensock: 3d foot reconstruction with flexible sensors," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020.
- [39] B. Xu, N. Wang, T. Chen, and M. Li, "Empirical evaluation of rectified activations in convolutional network," *arXiv preprint arXiv:1505.00853*, 2015.
- [40] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012.
- [41] A. Gretton, D. Sejdinovic, H. Strathmann, S. Balakrishnan, M. Pontil, K. Fukumizu, and B. K. Sriperumbudur, "Optimal kernel choice for large-scale two-sample tests," in *Advances in neural information processing systems*, 2012.
- [42] M. Long, Y. Cao, J. Wang, and M. I. Jordan, "Learning transferable features with deep adaptation networks," in *Proceedings of the 32nd International Conference on Machine Learning, ICML, 2015*.
- [43] I. Guyon, "A scaling law for the validation-set training-set size ratio," *AT&T Bell Laboratories*, vol. 1, no. 11, 1997.