



SeamPose: Repurposing Seams as Capacitive Sensors in a Shirt for Upper-Body Pose Tracking

Tianhong Catherine Yu
ty274@cornell.edu
Cornell University
Ithaca, New York, USA

Chi-Jung Lee
cl2358@cornell.edu
Cornell University
Ithaca, New York, USA

Ruidong Zhang
rz379@cornell.edu
Cornell University
Ithaca, New York, USA

Manru Mary Zhang*
mz479@cornell.edu
Cornell University
Ithaca, New York, USA

Cassidy Cheesman
crc268@cornell.edu
Cornell University
Ithaca, New York, USA

François Guimbretière
fvg3@cornell.edu
Cornell University
Ithaca, New York, USA

Peter He*
ph475@cornell.edu
Cornell University
Ithaca, New York, USA

Saif Mahmud
sm2446@cornell.edu
Cornell University
Ithaca, New York, USA

Cheng Zhang
chengzhang@cornell.edu
Cornell University
Ithaca, New York, USA




Figure 1: SeamPose repurposes seams as capacitive sensors in a shirt. Without modification to the clothing surface, the sensing shirt looks&wears similar to a conventional shirt and provides upper-body tracking capabilities. To make the sensing shirt, we machine-sew conductive threads over existing seams. Using 8 channels of capacitive seam signals (4 each on the left/right side) from the shirt, our customized deep-learning model estimates upper-body joint positions.

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UIST '24, October 13–16, 2024, Pittsburgh, PA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0628-8/24/10

<https://doi.org/10.1145/3654777.3676341>

ABSTRACT

Seams are areas of overlapping fabric formed by stitching two or more pieces of fabric together in the cut-and-sew apparel manufacturing process. In SeamPose, we repurposed seams as capacitive sensors in a shirt for continuous upper-body pose estimation. Compared to previous all-textile motion-capturing garments that place the electrodes on the clothing surface, our solution leverages existing seams inside of a shirt by machine-sewing insulated conductive threads over the seams. The unique invisibilities and placements of the seams afford the sensing shirt to look and wear similarly as a conventional shirt while providing exciting pose-tracking capabilities. To validate this approach, we implemented a proof-of-concept untethered shirt with 8 capacitive sensing seams. With a

12-participant user study, our customized deep-learning pipeline accurately estimates the relative (to the pelvis) upper-body 3D joint positions with a mean per joint position error (MPJPE) of 6.0 cm. SeamPose represents a step towards unobtrusive integration of smart clothing for everyday pose estimation.

ACM Reference Format:

Tianhong Catherine Yu, Manru Mary Zhang, Peter He, Chi-Jung Lee, Cassidy Cheesman, Saif Mahmud, Ruidong Zhang, François Guimbretière, and Cheng Zhang. 2024. SeamPose: Repurposing Seams as Capacitive Sensors in a Shirt for Upper-Body Pose Tracking. In *The 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*, October 13–16, 2024, Pittsburgh, PA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3654777.3676341>

1 INTRODUCTION

From the second-to-last Ice Age when humans put on clothes for warmth and protection [28], clothes have become indispensable to everyday life. Not long after Mark Weiser envisioned the future of computers to be woven “into the fabric of everyday life” [56], researchers explored smart clothing for wearable computing [19, 44, 44]. The always-on nature of clothing makes it an excellent medium for everyday pose tracking, a fundamental task with obvious applications in health care [40], human activity recognition [34], AR/VR interactions [1], human-robot interactions [26], sports analytics [27], etc.

Clothing deforms, stretches, and shifts as the joints and muscles move. Prior all-textile wearable movement sensing solutions have exhibited exciting performance by attaching conductive fabric patches across the clothing surface. More than 10 fabric patches covering different body parts were demonstrated to classify 10 upper-body movements [14] and to track continuous upper-body pose [63]. However, such modifications with patches of conductive fabric alter the base fabric’s properties: visual aesthetics and materiality (e.g., softness, stretchability, thickness, and breathability). As a result, the wearer’s experience changes, and the clothing designer needs to be aware of electrode placements’ impact on tracking performance [63].

To minimize surface modification and optimize the wearing experience while providing fine-grained tracking capabilities, we present SeamPose, which repurposes existing seams in a shirt as capacitive sensors for upper-body pose tracking. Seams are areas of overlapping fabric formed by joining two or more pieces of fabric together with stitches. Seam stitching is an essential step in the prevalent and scaled cut-and-sew manufacturing process that produces most everyday apparel. Seam placements are determined by the pattern, a set of templates designed for cutting and sewing the fabric into garment [47]. Because seams originally exist on the garment and remain concealed when worn, altering seams with conductive threads will not change the appearance or the materiality of the apparel, while providing exciting tracking capabilities.

While researchers have extensively explored with conductive or functional stitches [4, 13, 18, 21, 42, 61], as we discussed, large areas of conductive patches are needed for the complex body pose estimation task in the past. This paper aims to answer the research question: **whether capacitive sensing seams, repurposed from existing seams, on a shirt can estimate upper-body pose.**

To answer this research question, we developed a proof-of-concept prototype based on a common and basic long-sleeve shirt. To transform the seams into capacitive sensors, we machine-sew insulated conductive threads over the existing seams. It is important to note that we only augment existing seams from the selected shirt pattern and do not strategically add electrodes to locations that better capture body movements [7, 15, 32, 47]. This prototype is untethered and battery-powered, as shown in Fig. 4. There are 8 seam electrodes (Fig. 5), 4 each on the left and right side of the shirt. These eight seam electrodes are connected to a customized active capacitive sensing board that measures and transmits the signals wirelessly via Bluetooth. The sensing principle of our approach is that *different body poses and movements will deform seam electrodes and change the coupling between the human body and seam electrodes, leading to unique and complex patterns in measured capacitances.*

To extract and interpret the pose information from these complex capacitance readings, we customized a deep-learning pipeline that estimates 8 relative (to the pelvis) upper-body joint positions in 3D. To evaluate SeamPose, we conducted a user study with 12 participants. The results showed that SeamPose tracks upper body poses with a mean per joint position error (MPJPE) of 6.0 cm, comparable with prior wearable pose tracking systems. In summary, the main contributions of this paper are:

- We described a fabrication process to repurpose existing seams in a long-sleeve shirt into capacitive sensors for body pose tracking, one step towards minimally obtrusive conductive textile sensing.
- We developed a proof-of-concept untethered long-sleeve shirt prototype and a deep learning framework that estimates upper body joint 3D positions from the capacitance measured by these conductive seams.
- We conducted a user study with 12 participants and achieved promising results, as a proof-of-concept to verify the feasibility of this proposed approach.
- We further discussed the opportunities and challenges of generalizing SeamPose on various clothing types and widespread adoption in everyday life.

2 RELATED WORK

SeamPose tackles pose tracking with capacitive sensing seams in a long-sleeve shirt. Capacitive sensing with textiles exhibited exciting potentials in strain sensing [6], environmental sensing [2], object recognition [57], gesture recognition [4, 24, 45], etc. Discussing the fundamentals of capacitive sensing is beyond the scope of this paper, we refer readers to Grosse-Puppenthal et al. and Bian et al. for comprehensive reviews of body-area capacitive in HCI [8, 20]. In this section, we will focus on discussing the prior work that is closely related to wearable pose-tracking including 1) wearable pose estimation, and 2) body movement sensing with smart clothing.

2.1 Pose Estimation with Wearable Sensors

Compared to high-fidelity vision-based MoCap systems [16, 43, 53], wearable solutions enable on-the-go pose estimation without the need for setups in the environments. Xsens [41] attaches 17 wireless inertial measurement units (IMUs) onto a tight-fitting suit to enable

professional-grade full-body captures. Decreasing the number of instrumentation sites alleviates the cumbersome setup but drastically complicates the tracking task due to limited sensed information. Researchers explored full-body tracking with six IMUs [22, 58], six electromagnetic sensors [25], or four flex sensors [11] attached onto tight fitting suits. However, tight-fitting suits are not comfortable for everyday uses. In AR/VR uses [38], full-body poses can be inverse-kinematically inferred with sensors in the headset and hand controllers [3, 5, 23]. IMUPoser uses IMUs in consumer smartphones, smartwatches, and earphones to infer full-body pose [39]. A single camera could track body pose when mounted on the head [52, 54]/wrist [33] from egocentric/partial body views, respectively. Recently, smartglasses have utilized ultrasonic sensing to track upper-body poses [37]. SeamPose shares the approach of pose reconstruction with limited sensed information for minimally obtrusive integration and contributes a new approach with seam sensors in clothing that afford everyday uses.

2.2 Body Movement Sensing with Smart Clothing

Clothing covers a large area of human bodies and the always-on nature of clothing affords always-on body movement sensing. Attaching distributed miniaturized electronics (e.g., IMUs [22] and flex sensors [11]) to the garments hinders the softness and requires either wireless sensor network [41] or optimized on-body wirings [55].

In contrast, all-textile systems, commonly employing conductive fabrics/threads, augment garments with motion-sensing capabilities while preserving the all-textile softness. Liang et al. bonded conductive fabric patches onto tight-fitting leotards to coarsely monitor dance movements with resistive changes [32]. Esfahani et al. sewed 11 polymerized resistive patches on an undershirt to classify 10 upper body movements [14]. Most resistive strain-based approaches require tight-fitting garments, similar to that of IMU suits, because the sensors must be firmly attached to the expected body locations [65]. Gioberto et al. used resistive coverstitch on loose-fitting garments to detect fabric bends and folds [15]. MoCapaci integrated 4 textile cables as capacitive antennas into a blazer and effectively classified 20 upper-body poses [7]. In contrast to works above that analyze the correlation between signals and movements or perform classification tasks, SeamPose aims to infer upper-body joints in 3D.

The most recent work, Mocapose [63], is the only existing loose-fitting all-textile system that accomplished the upper-body 3D joint reconstruction task. With capacitive signals from 16 channels (8 each on the left/right) of conductive fabric patches glued onto a jacket, a 2D CNN-based deep learning model demonstrated strong performance in tracking upper body pose. In contrast, our proposed approach alters the garment at the thread-level instead of the at the fabric-level, a step towards minimizing the alteration of the garment visually and tactilely. By only sewing over the 8 existing seams with conductive threads, SeamPose can track the upper body pose continuously on a loose-fit long-sleeve shirt in real-time.




Figure 2: The seams and patterns of our proof-of-concept prototype with a long-sleeve T-shirt. The black fabric pieces represent the patterns of the T-shirt, while the dotted lines are the repurposed seams, also indicating where the stitches are when joining the fabric pieces. The color of the dotted lines can be mapped to the lines, or seams, on the constructed T-shirt.

3 SEAMS & PATTERNS

For garments, *patterns* serve as templates representing the shapes and sizes of the fabric pieces required to construct a specific garment. *Seams*, on the other hand, are the areas where the fabric pieces overlap when sewn together to make the garment (Fig. 2). Consequently, patterns define the number of seams, seam placements, and the shapes of fabric to be seamed. Patterns affect the garments’s structure, size, fit, etc. Although seams vary in their placements, quantity, and length, they are generally distributed throughout the garment. This distribution provides opportunities to collect rich information regarding the body. As a result, in this paper, we leverage this feature and repurpose the seams as capacitive sensors to track body poses.

Because we are repurposing existing seams as sensors, patterns further decide the sensors’:

- **Placements:** jackets, blazers, and collard shirts commonly have seams in the front and back of the torso for fit, but T-shirts, sweaters, and sweatshirts generally do not;
- **Quantity:** jackets could easily have twice as many seams as T-shirts, and even for a long-sleeve T-shirt, some patterns have 8 seams¹ while others² have 10 seams (2 additional ones on the sides of the torso).
- **Lengths:** the sleeve seam in a long-sleeve shirt covers the elbow, an important joint to track, while a short-sleeve shirt does not.

For the proof-of-concept prototype, we chose a long-sleeve T-shirt³. Long-sleeve T-shirts are basic and common. They have the fewest numbers of seams among cut-and-sew long-sleeve tops. Our selected shirt has 5 seams on each side (Fig. 2): 1 above the shoulder, about 11cm long, 1 in front of the shoulder, about 30cm long, 1 behind the shoulder, about 30cm long, 1 along the sleeve, about 49cm long, and one on the side of the torso, about 44cm long. To ensure minimal alteration and maximal generalizability, we chose not to repurpose the ones on the sides of the torso because not all long-sleeve shirt patterns have those seams, as explained above.

¹<https://a.co/d/6mg3bW8>

²<https://www.uniqlo.com/us/en/products/E460354-000>

³<https://www.michaels.com/product/long-sleeve-crew-neck-adult-t-shirt-by-make-market-M20033775>

Although our trained models do not aim to generalize across different patterns, adding more seam electrodes will provide additional information that likely will further improve the tracking performance. The purpose of this paper is to demonstrate the feasibility of this approach and present a baseline for future exploration. In addition, we acknowledge that the results presented in our paper can not be directly replicated on sleeveless garments like strapped and strapless tops.

4 SEAMPOSE IMPLEMENTATION

We repurpose seams as capacitive sensors for upper-body pose tracking, without modifying the clothing surface. SeamPose prototype has three main components:

- A shirt with 8 conductive seam electrodes (detailed in Sec. 4.1): the electrodes are symmetric on the left/right side of the shirt with 1 above the shoulder, 1 in front of the shoulder, 1 behind the shoulder, and 1 along the sleeve;
- A customized capacitive sensing board, detailed in Sec. 4.2; and
- A deep learning pipeline estimating the positions of 8 upper body joints from the readings of capacitive sensors (Sec. 4.3), detailed in Sec. 4.4.

4.1 Conductive Seams Fabrication

To transform conventional seams into capacitive sensing seams, we machine-sew conductive thread over existing seams, as shown in Fig. 3. The base fabric of our selected unisex shirt (Make Market Long Sleeve Crew Neck Adult T-Shirt) of size medium is 100% single knit cotton jersey. We use a home sewing machine (SINGER Heavy Duty 4423 Sewing Machine) with a sewing needle of size 80/12 (SCHMETZ Universal 130/705). The top thread uses a conventional polyester sewing thread. Similar to prior works that use functional threads/wires as bobbin threads to relieve mechanical stress [4, 42, 61], the bottom bobbin thread uses an off-the-shelf TPU-coated 2-ply silver-plated nylon thread (Shieldex 117/17 x2 HCB TPU, Fig. 3(B)) with a resistance profile of $< 300\Omega/m$. We choose a zigzag stitch, common for seaming with home sewing machines. After extensive experiments for a clean and consistent finish, we set the thread tension, stitch length, and stitch width to 5, 2.5, and 2.5, respectively.

Conductive Thread Selection. The choice of conductive thread can heavily impact the overall system performance. In our early prototypes, we experimented with a non-insulated silver-plated nylon thread (LessEMF, $< 100\text{ Ohm/cm}$). We noticed (a) electrode shorting caused by body movements so we looked into insulated conductive threads that are compatible with home sewing machines, and (b) frequent conductive thread breakage when sewing. For our selected TPU-coated thread, the elastic TPU coating (a) insulates the conductive core, and (b) provides strong mechanical properties allowing $21\% \pm 5$ elongation before breaking⁴. When prototyping with the TPU-coated thread, the conductive thread never broke when sewing. The strong mechanical properties not only increases the fabrication




Figure 3: Machine Sewing Conductive Thread. (A) Home sewing sewing machine setup. (B) The off-the-shelf insulated conductive thread in use. (C) An illustration of machine-sewn conductive thread traces over existing seams, overlapping areas formed when stitching (by the original seam stitches) pieces of fabric together. (D) A side-view illustration of how the TPU-insulated conductive thread with silver nylon core is stitched onto the fabric.




Figure 4: The battery-powered customized sensing board is housed inside a 3D-printed PLA case and hot-glued onto the prototype below the neck.

success rate but also strengthens the prototype’s durability, important for wearable applications with frequent re-wearing. Our final prototype has been tested with over 20 people and hundreds of re-mounting without breaking.

Connectors. Electrically and mechanically robust connections between textile-based sensors and rigid electronics remain an open research problem [51]. We connect our insulated threads similar to insulated wires with DuPont wire-to-wire connectors. We crimp the stripped conductive thread core inside the conductor tab and the insulating TPU-coated thread inside the insulation tab. This method ensures stable connections during movements and re-wearings.

⁴<https://www.shieldex.de/wp-content/uploads/2021/05/Y-VTT-Datasheet-Shieldex-117-17-x2-HCB-TPU-V4.pdf>




Figure 5: Example SeamPose Signals. We show 19 seconds of continuous SeamPose the left and right seam signals (4 each) as the wearer performs symmetric (both arm raises) arm movements, asymmetric arm movements (right arm curls, left front raises, alternate punches, confused gesture), head movements, and standing still. The colors map the signals and their corresponding electrode placements, illustrated on the left. The numerical value of signal on y -axis is median normalized.

4.2 Customized Sensing Board

The purpose of the customized sensing board (Fig. 4, 36x31mm) is to measure and transmit the capacitance of the connected conductive seams without hindering body movements. We replicate the single-ended configuration circuit in FDC2214 (Texas Instruments) evaluation module⁵: the 4-channel capacitance-to-digital converter measures capacitances with great resistance to electromagnetic interference. We chain 2 FDC2214s on I2C to read 8 channels, at a configured sample rate of 32Hz. XIAO nRF52840's (Seeed Studio) transmits the sensor readings via onboard Bluetooth Low Energy (BLE) functions to a nearby computer. We use Arduino to program the firmware and UART for Bluetooth communication. We power the circuit with a 3.7V 290mAh Lipo battery. The mean measured⁶ power consumption is 12.7mA, which can continuously transmit measurements for 22.83h. Approximately, the board weighs 11g and costs US\$27.

4.3 SeamPose Signals

Changes in the wearer's upper body pose cause changes in the seam electrode's sensed capacitances, which are input into our deep learning pipeline that infers the relationship between the body pose and measurement capacitances. When connected to the sensing board, the conductive seams become self-capacitance sensor electrodes. In a loading-mode capacitive sensing system, the setup is simple: each sensor electrode acts as both a transmitter and a receiver [50]. In any loading-mode system, electrical current displacement occurs as the electrode's capacitive coupling with the surrounding changes. In SeamPose, the seam electrodes are coupled with the wearer's body, so when the body pose changes (*i.e.*, limb movements), the seams move along with the limbs, leading to changes in the coupling between the body and seam electrodes, hence the current

displacement changes. In addition, the seam electrode itself deforms, distorts, and displaces during body movements, altering the electrode's capacitance which also affects the current displacement. In a nutshell, we conjecture the capacitive changes are attributed to (1) coupling changes between the seam electrodes and the wearer's body, and (2) seam electrode's self-capacitance changes.

Notice that in Fig. 5, the sleeve electrodes (purple) not only trace along the sleeve seams but also trace along the seams in the front (beige) of and on top of the shoulder (blue). This is a design choice to avoid additional wire routings across the shirt: one ends of all electrodes meet at the sensing board. Even though some sensor placements "overlap", they still individually provide sensing information, which we further validate in Sec. 6.5.

As shown in Fig. 5, the seam signals correlate to the wearer's movements. The sleeve electrodes, colored purple, have the largest changes in magnitude as they are the longest electrodes with the largest range of motion. The shoulder front (colored beige) and back (colored green) electrodes have similar magnitudes as they are symmetrically placed, but the relationship between these two signals depends on the motion. For example, for the both-arm side raises, the two signals both increase as the arms raise and decrease as the arms return, but for alternate punching movements, one signal increases as the other decreases. The shoulder top (colored blue) has the smallest magnitude for its shortest length and relatively stationary nature, compared with other electrodes that touch moving joints. Though the magnitude is small, it is proven to still be information-rich in our later evaluation.

For asymmetric movements (*e.g.*, right arms curls and left front raises), the static side still exhibits signal changes with small magnitudes. Moving one side of the shirt often stretches the other side because different parts of the shirt are interconnected and body movements involve complex muscle engagements. For example, a one-arm raise could (a) stretch the back of the shirt which changes the signals as the sensing board is in the back, and (b) alter the whole

⁵<https://www.ti.com/tool/FDC2214EVM>

⁶<https://lowpowerlab.com/guide/currentranger/>




Figure 6: Customized Deep Learning Pipeline.

body’s coupling with the shirt. Further, at the thread-to-sensing-board connection, some insulated threads are close to and even overlap with each other and thus attributed to crosstalks, which is a limitation to be addressed in the future, but in our prototype, we input the 8 seam signals together, instead of side-by-side, to avoid possible confusion to the DL model.

Although our prototype does not cover or directly instrument the head, it can sense the head movement (see head roll signals differing from standing still signals in Fig. 5) based on overall body-capacitance changes and subtle shoulder movements. To extract and interpret the pose information from these complex capacitance changes, we developed a data-driven approach.

4.4 Deep Learning Pipeline

With eight channels of capacitive signals from the seams on the shirt, we estimate the wearer’s upper body pose in 3D, relative to the pelvis. In the previous subsection, we show the correlation between body movements and seam signals, but the complex mapping between the 8 channels and the body pose is not immediately clear. The reconstruction task is challenging because the tracked body parts have a total of 19 degrees of freedom [9], and we only have sparsely 8 channels of 1-dimensional temporal signals, constrained by our design choice that minimally alters everyday clothing. Unlike tight-fitting suits, the sensors are not firmly fixed to a mapped body part, complicating the problem [65]. The task is further complicated by the soft cotton fabric having natural draping variations. For example, when the wearers raise their arms over the head, the sleeves naturally slide down the arms and attribute signal variations. Embracing these technical hurdles in pursuit of comfortable and seamless pose-tracking integration into everyday clothing, we customized a deep-learning pipeline to learn the complex mapping between the 8 input channels.

4.4.1 Ground Truth Acquisition. Recent computer vision advancements in pose reconstruction with a single RGB image have proven more accurate for users in loose-fitting cloth [46]. Similar to prior works on non-vision-based pose tracking [10, 29–31, 60, 63], we selected state-of-the-art computer vision models for ground truth acquisition. Given an RGB image frame, Detectron2⁷ first detects the area containing the person, then HMR2.0 [16], a vision transformer-based model, estimates the SMPL [35] representation of the detected

person. SMPL is a mathematical model that describes a human body mesh template of 6890 vertices with shape components β ($\dim(\beta)=10$) and pose components θ ($\dim(\theta)=216 = 24 \text{ joints} \times (3 \times 3) \text{ rotation matrix}$ representing the rotation from its kinematic prior). Among the 24 joint components, the pelvis provides global orientation (*i.e.*, the direction the body is facing) and the other 23 joint components represent rotations from their kinematic prior. SMPL also acts as a forward-kinematic model and calculates joint positions, J , based on joint rotations, θ . Note, SeamPose does not predict body shape or global orientation. The ground truth labels have 2 folds:

- 13 pose components of size $[13 \times 6]$: rotations of 13 upper body joints, and we exclude the hands, similar to [5, 23]. SMPL uses 3×3 rotation matrices to represent rotations, but a 6D rotation representation has been proven effective [23, 39, 64], so we convert rotation matrices into 6D rotation vectors.
- 8 joint positions of size $[8 \times 3]$: 3D position representation of upper-body joints, in x,y,z coordinates, calculated by SMPL with pose components: nose, neck, right shoulder, right elbow, right wrist, left shoulder, left elbow, and left wrist. We scale the joint positions into physical units (meters) with measured arm lengths and center the pelvis at the origin.

4.4.2 Input Normalization. The input to the model is a 3s-window (96 frames) of 8 channels of seam signals. The capacitive readings depend on the body pose and the capacitance of the seam electrodes. Because the electrodes are not of equal length and are distributed on different parts of clothing, channels of signals are not centered together. We first calculate the median of the most recent window of 5.6s (180 frames) and perform a median normalization. And to better center individual channels around 0, we subtract each channel with: $0.98 \times \text{channel median}$ within the 3s-window. The scaling factor aims to preserve inter-channel relationships. For example, without the scaling factor, standing still for more than 5.6s with arms down and arms up will have the same normalized signals.

4.4.3 Model. The model architecture is detailed in Fig. 6. The input signals, of dimension $[96 \times 8]$, are first transformed into an embedding of dimension 96 with a linear layer. Then, the embeddings go through a 2-layer bidirectional LSTM (inspired by [11, 22, 39, 66])

⁷<https://github.com/facebookresearch/detectron2>

of hidden dimension 256. LSTM is chosen to better model the temporal patterns, but there are potential optimizations on model architectures (e.g., transformers) for future works. We take the last 3 prediction frames (0.09s) and decode the embeddings with 2 linear layers with out-feature sizes of 256 and 78 (for the 13 pose components). We then MedianPool the last 3 prediction frames to mitigate jitter [11]. Finally, we use SMPL to calculate joint positions based on the pose components. Our model directly regresses the pose components representing joint angles that are more transferrable across body sizes. The conversion to joint positions allows easy comparisons with other systems as the position error is the common metric for wearable pose tracking. We further smooth the joint position predictions with a running median filter of window size 5 (0.17s).

4.4.4 Training. We implemented the models in PyTorch and trained them on an NVIDIA GeForce RTX 2080 Ti. We use the Adam Optimizer with a cosine learning rate scheduler. To compute the loss, we add the mean absolute error (MAE) losses of pose components and joint positions:

$$\mathcal{L} = \|\theta - \theta^*\|_1 + \|J - J^*\|_1 \quad (1)$$

Here, θ and J are the predicted 13 pose components and 8 joint positions. θ^* and J^* are the ground truth 13 pose components and 8 joint positions. The batch sizes are 512.

We adopt a two-stage training scheme: (1) User-Independent training stage: we first train a user-independent model with 15 epochs and a starting learning rate of $8e^{-3}$; and (2) User-Adaptive training stage: we then finetune the trained user-independent model with another 10 epochs and a starting learning rate of $4e^{-4}$.

4.4.5 Data Augmentation. To improve the model’s robustness against different body sizes and movement patterns, we apply data augmentation techniques to introduce variations to the training data in each epoch. To mitigate window and channel median variations that impact the input normalization, at 80% chance, we apply random shifts to window and channel medians independently. To account for signal range variations caused by the fit of the shirt, at 80% chance, we apply random scaling in the range of [94%, 106%] to the normalized input. To further introduce randomness, at 80% chance, we scale individual reading in the range of [99.7%, 100.3%].

4.5 Real-Time Inference Pipeline

We implement a real-time end-to-end inference pipeline (Fig 7). For real-time joint visualization, we use the visualizer provided by EasyMocap [49]. On an Apple Macbook Air (2022), the inference (10.7ms) and visualization (16.0ms) latency are 28ms in total.

5 USER STUDY EVALUATION

To evaluate SeamPose’s continuous upper-body pose estimation performance, we conducted a user study, approved by the Institutional Review Board (IRB). We recruited 12 participants (6 self-identified as male, 6 as female, mean age=24.9, std age=4.0) spanning a variety of body shapes, detailed in Table 2. Each study lasted about 1 hour and compensated US\$15.

We conducted the study in an experiment room on a university campus. At the beginning of the study, the experimenter instructed




Figure 7: Real-Time Inference with SeamPose. Note, for our MPJPE calculation, the head is represented by only the nose key point to avoid error dilution, but for visualization purposes, we include eyes and ears key points reconstructed by SMPL.

the participant to stand in front of a green screen. A laptop (Apple Macbook Air, 2022) was placed on a desk about 3m away from the participant. The laptop continuously recorded (a) ground truth video via its built-in camera (30fps) and (b) SeamPose sensing data received via Bluetooth from the prototype described in Sec. 4 and synchronized the two with its clock. The laptop and its connected monitor displayed visual stimulus for movements and the camera view that captures the full body. We chose video stimuli over photos and text as videos contain details of the movements. Participants were asked to follow the movements on the video but were not strictly asked to follow the same pace nor the exact movement patterns (e.g., golf&tennis swings and dance movements varied greatly among participants).

For each participant, we collected 8 sessions of data. Before each session, the participant was instructed to take off and put back on the shirt themselves in order to evaluate our system across different wearing sessions. Participants wore their own clothes of various types and sizes that fit underneath our prototype shirt. Each session lasted about 227 seconds. Each session contains 3 sections, including three types of upper-body movements informed by prior work:

- **Section 1: 54 randomized unique movement videos (195s)**, detailed in Appendix A, cover casual/daily gestures [7], sports movements [5, 11], and controlled terminal poses [37] movement sequences [63] which explore poses that are uncommon in daily activities but kinematically feasible.
- **Section 2: 2 TikTok dance videos^{8,9} (21s)** introduce fluent movements and rare poses [63].
- **Section 3: Freestyle movements (10s)** introduce unseen poses/movements that are not included in our defined set. Participants are instructed to perform random movements of their choice, not limited to the upper body.

As a proof-of-concept on a research prototype, we can not exhaust all possible upper-body poses. The three different types of movements are chosen to provide enough variance of the body poses

⁸<https://www.tiktok.com/@sophielaverie/video/7334053266050911521>

⁹<https://www.tiktok.com/@sophielaverie/video/7326633661162474784>




Figure 8: Joint Error Breakdown. In (A), we show the distance error distributions of the 8 predicted joints, labeled in (B). Error bars represent the standard deviation across participants.

to demonstrate the potential of this proposed sensing approach to tracking body poses.

In total, we collected 1813.6 seconds of data from each participant, resulting in $1813.6s \times 30\text{fps}$ (camera’s sampling rate) = 54408 training/testing instances. From all 12 participants combined, we collected 6.03 hours of data.

At the end of the study, participants completed a questionnaire collecting information about their demographic, body sizes (measured by the experimenter), and the prototype’s wearability.

6 RESULTS

We evaluated SeamPose for both user-adaptive and -independent scenarios, as well as estimating performance for each body joint and various motions. Additionally, we conducted an ablation study to understand the impact of seam placements on SeamPose’s continuous tracking performance.

Evaluation Metrics. Informed by other wearable pose tracking systems [5, 33, 37, 63], we adopt Mean Per Joint Position Error (MPJPE) as the evaluation metric for continuous pose tracking: the mean Euclidean distance errors of 8 joint positions in centimeters (cm), relative to the pelvis.

6.1 User-Adaptive Model Results

As mentioned in Sec. 4.4.4, we first trained a user-independent model, without the participant’s data in the training set, and then we fine-tuned the user-adaptive (UA) model with the the evaluated participant’s data. To simulate the user calibrating the device before they start using it, we use the first 6 sessions (22.7 min) of data from the evaluated participant for fine-tuning and testing on the last 2 sessions. Note that the participant took off and put back on the shirt before every session, our system is session-independent. SeamPose achieves an overall MPJPE of 6.0cm (std=0.65cm). We recognize that providing 22.7 min of calibration data may not be always preferred for the optimal user experience. Therefore, we conducted further experiments exploring how much fine-tuning data is required to achieve a good tracking performance. As shown in Fig. 9, reducing the fine-tuning data to 18.9min, 15.1min, 11.3min, 7.6min, and 3.8min, MPJPE increases to 6.1cm, 6.2cm, 6.4cm, 6.6cm, and 7.0cm,




Figure 9: As the user provides more training data for the user-adaptive model, the model’s tracking performance increases. Error bars represent the standard deviation across participants.

respectively. With only 3.8 minutes of user-specific training data that contains only 1 repetition of each movement, SeamPose still has a promising tracking performance of 7.0cm (std=0.68cm), 1cm worse than that with 22.7min of training data.

6.2 Unseen User: User-Independent Model Results

For an even better user experience, the shirt should be able to track body poses "out-of-box" without any calibration data from the new user, *i.e.*, user-specific training data. We adopt a leave-one-participant (LOPO) cross-validation to evaluate our performance in this scenario. User-independent (UI) models were trained 7.75 hours of data: 30.2 minutes of data from each of the other 11 participants and 5 hours of data from 7 researchers. As shown in Fig. 9, the MPJPE increases to 8.6cm (std=0.93cm), 1.6cm worse than 3.8min of training data, and 2.6cm worse than 22.7min of training data. The performance degradation is expected because individuals’ body sizes and capacitances vary, but SeamPose still yields low-fidelity




Figure 10: Movement Error Breakdown. Error bars represent the standard deviation across participants.

tracking, sufficient for certain applications like activity recognition and lifelogging.

6.3 Results Analysis based on Joints

To help better understand our performance, we further analyzed our performance based on individual joints, as illustrated in Fig. 8. Similar to prior work [5, 39, 63], the wrists with the most moving distance have the largest errors: 12.2cm, 14.4cm, and 16.4cm with 22.7min, 3.8min, and 0min of user data respectively. The end effectors accumulate errors along their long kinematic chains and the wrists have large ranges of movements. The elbows have the second largest errors: 6.7cm, 7.8cm, and 9.1 cm with 22.7min, 3.8min, and 0min of user data. The neck and shoulders have the smallest errors, as they did not have much movement compared to wrists and elbows. We notice a big difference in estimating the position of the head between the UI model (6.4cm) and UA model (3.8cm) with 3.78min of data. We conjecture this is caused by the fact that SeamPose does not directly instrument the head, unlike other tracked joints that affect the seam electrode shapes, so head movements are mostly inferred from the overall coupling change between the body and the seams. This indicated that our system may need calibration data for a reliable estimate of head position.

6.4 Results Analysis based on Motion

We further break down the results into motion types. Like many other data-driven wearable MoCap systems, the performance decreased for estimating unseen movements. As we discussed in Sec. 5, we cannot exhaust all possible upper-body movements&poses for evaluation. However, when we designed the study, we explicitly included dance and freestyle movements to better understand our system’s limitations on extreme and unseen poses and movement patterns. As shown in Fig. 10, across all three models, defined movements (5.9cm, 6.8cm, and 8.3cm with 22.7min, 3.8min, and 0min of user data) have the smallest errors, followed by dance movements (8.2cm, 9.5cm, and 10.8cm with 22.7min, 3.8min, and 0min of user data) and freestyle movements (10.8cm, 12.1cm, and 13.5cm with 22.7min, 3.8min, and 0min of user data). Adding user-specific training data also consistently improves the tracking accuracy for each movement type. This result suggested that the system similar to many data-driven systems, can further benefit from being




Figure 11: Seam Removal Ablation Study Results. A larger MPJPE indicates that the removed seam is more important. Error bars represent the standard deviation across participants.

trained with a much larger dataset containing a large variety of body movements. However, as a proof of concept, it still showed reliable tracking performance on unseen movements and subjects.

6.5 Seam Removal Ablation Study

For our prototype, we chose a long-sleeve shirt with 8 seams, 4 each on the left and right side of the shirt. Intrigued by the question of how different seam placements contribute to the model and how our approach will generalize with even fewer seams on a short-sleeve shirt and sleeveless shirt, we conducted a seam removal ablation study on the 6 odd-numbered participants. Because the seam placements are symmetric, we remove 2 symmetric seams (e.g., left and right sleeve electrodes together) at a time, and train with the rest 6 seams. We trained user-adaptive models and user-independent models for the 4 possible seam removals. As shown in Fig. 11, with no surprise, given the small number of signal channels for a complex tracking task, removing any seam in any model leads to an increased error. However, some seam placements have larger effects than others.

As we discussed in Sec. 4.3, the shoulder top seams have the smallest changes in magnitude, removing them does not harm the performance greatly: 0.2cm, 0.2cm, and 0.5cm worse with 22.7min, 3.8min, and 0min of user data. The shoulder front/back seams have similar significance across all models, and removing them increases the errors by 1.0cm/0.9cm, 0.8cm/1.0cm, and 1.0cm/1.1cm with 22.7min, 3.8min, and 0min of user data. On the other hand, the sleeve electrodes’ importance varies across models. For both adaptive models, the sleeve electrodes have the most important placements (*i.e.*, their removal causes the largest errors compared with other placements) which lead to about 1cm additional error when removed. For the user-independent model, the sleeve electrodes become the least important placement and only lead to an additional error of 0.3cm. The observation implies that signals from some placements (sleeve) generalize across users worse than others do. Furthermore, given the reasonable performance across all three models with the sleeve electrodes removed, SeamPose has the potential to generalize on short-sleeve shirts and even sleeveless ones.

Table 1: Situating SeamPose’s tracking performance in the literature. “UI” stands for user-independent MPJPE. “UD” stands for user-dependent/adaptive MPJPE.

| | Tracking Device(s) | UI (cm) | UD (cm) |
|------------------------|--|---------|---------|
| MI-Poser [5] 2023 | magnetic and inertial sensor fusion on a headset&2 controllers | 6.6 | - |
| PoseSonic [37] 2023 | acoustic sensing on smartglasses | 6.2 | 5.6 |
| LIP [66] 2024 | 4 IMUs on a jacket | 10.58 | - |
| MocaPose [63] 2023 | 16 conductive patches on a jacket | 8.8 | 8.6 |
| SeamPose | 8 conductive seams on a shirt | 8.6 | 6.0 |

6.6 Performance Comparison with Prior Works

SeamPose is the first to repurpose seams into capacitive sensors in clothing for body pose tracking. Direct comparisons with other works are difficult because SeamPose and some other related works used different datasets to evaluate the performance. Factors like training data duration, pose/movement set design, data collection setup, etc. heavily affect the performance. Thus, we only include the comparison for an empirical comparison to help readers to situate our performance in the literature. As this table 1 showed, SeamPose performance is comparable with prior wearable pose tracking systems while keeping the appearance of the shirt largely unchanged.

6.7 Perceived Comfort

In addition to its promising tracking ability, our sensing approach aims to preserve the soft and comfortable nature of clothing. Based on the survey results, the participants felt very comfortable (Median=5 on the 5-point Likert scale; 1=very uncomfortable, 5=very comfortable) with the shirt prototype and perceived it as similar to their everyday clothing (Median=5 on the 5-point Likert scale; 1=very different, 5=not different at all).

7 DISCUSSION, LIMITATIONS & FUTURE WORK

SeamPose demonstrates a minimally obtrusive continuous upper-body pose-tracking solution powered by a minimally altered shirt. The proof-of-concept prototype system was evaluated in a lab study to demonstrate its feasibility in tracking upper body pose. In this section, we discuss the limitations of our current implementation and the opportunities and challenges of ubiquitous seam-enabled pose-tracking integration into everyday apparel.

7.1 Improving Tracking Performance

From prototyping the proof-of-concept shirt, we found the seam signal quality is mainly influenced by (a) the thread selection: conductivity and insulation (Sec. 4.1); and (b) the positions and number of seam electrodes (Sec. 3). We chose to constrain the latter for our design choice of minimal alteration and maximal generalizability and explored the positions and the number of seams with the seam

removal ablation study, but an extensive characterization of such factors may improve the information gain.

Although our tracking performance is comparable with state-of-the-art loose clothing pose-tracking systems, we observed the drop in performance with unseen users and unseen movements (detailed in Sec. 6), a limitation that plagues many data-driven wearable sensing technologies. Inspired by recent IMU-based pose-tracking advancements empowered by large synthetic datasets (generated from attaching virtual IMUs to SMPL meshes) [22, 39, 66], we plan to explore simulating seam electrodes’ capacitive measurements to address this issue. Prior works have simulated capacitive measurement on the human body for activity recognition [12], and more recently, Schöffmann used a 3D finite-element method to simulate capacitive touch readings for human-robot interaction [48]. Additionally, other thread-based sensing approaches like triboelectric, piezoelectric, and impedance sensing [36, 59, 62] can be integrated into seams and become multi-model seam sensors to complement the current system and further improve the tracking performance.

As shown in Fig. 9, adding more user-specific data to the current user-adaptive model seems to not significantly improve the performance. In the future, we plan to make prototypes of different sizes (*i.e.*, small/medium/large) and introduce user descriptors (*e.g.*, anthropometric data) as inputs to the deep learning pipeline to better learn the variations among different body shapes. Moreover, more data from other users at a significantly larger scale for training the user-independent model may improve both the user-independent and -adaptive performances as shown in many prior systems.

Another limitation of our current implementation is that we do not estimate global orientation, even though relative joint tracking suffices for applications like activity recognition, rehabilitation, etc. For use cases that require global orientation, SeamPose could become a part of a multi-device pose tracking framework (*e.g.*, along with headsets for hands-free AR/VR uses or smartphones for mobile interactions).

7.2 Manufacturing at Scale for Everyday Wearing Experience

In this paper, we repurposed the seams by machine-sewing conductive threads over existing seams. The fabrication process is simple and holds tremendous potential for integration into industrial processes. If the insulated conductive thread (Sec. 4.1) can withstand mechanical stress induced by industrial sewing machines and/or sergers, not explored in this paper, the sensing seams can be easily integrated into the industrial scaled cut-and-sew process by simply replacing the bobbin thread with conductive thread. Instead of repurposing existing seams like the method we propose in this paper, seams can be functional as soon as they are created at factories.

Beyond industrial mass manufacture, we share some challenges with current research related to garments when aiming for everyday uses, *i.e.*, the garments need to be waterproof and wear-proof. One key avenue is in designing a washable sensing shirt with a detachable or washable sensing board [51]. Further, extensive characterization of durability and possible performance degradation over wears (*e.g.*, hysteresis), and mitigation approaches are essential for real-world adoption.

7.3 Evaluation in Real-World Settings

We evaluated SeamPose extensively in a controlled lab setting as a proof of concept to justify the feasibility of this approach. However, there are scenarios that our system needs to be further tested for real-world deployments. For instance, while the participants wore their own shirts (different materials with different dielectric constants and long-/short-/no sleeves with different amounts of skin contacts) underneath the shirt, we did not investigate putting clothes over our prototype, which is common when wearing a shirt.

Moreover, our seam electrodes are capacitive sensors that inherently sense touches and proximity. Because the seam electrodes are very small in size, they are very insensitive to changes in even near proximity [17, 63]. At the signal level, there are no changes in the seam signals when the wearer touches electrically grounded and other objects that could cause electromagnetic (EM) interferences, but there are small (relative to body movements and direct electrode contacts) changes when other objects/people approach the electrodes within about 5cm. This finding is consistent with prior work [63]. As the tracking level, with our real-time pipeline, there is no obvious performance degradation in the above situations. However, it remains rather sensitive to touches and direct contact, especially heavy ones causing thread deformation, which may happen during intense exercise or while sleeping. Therefore, to deploy the system for everyday pose-tracking wearing experiences, the system needs to be fine-tuned to be robust to these noises. We plan to take a data-driven approach to collect more data in these scenarios and update our machine-learning model to handle the noise as mentioned above.

SeamPose focuses on leveraging seams in a shirt for upper body motion estimation. Similar to prior work [63], we were only able to evaluate our approach in one shirt with one kind of commonly seen seam placements. However, there are many different clothing patterns. For example, seam electrodes around the wrists at the sleeve cuffs could improve/enable forearm/palm orientation tracking, respectively. The seam pattern we chose has one of the smallest numbers of seams for long-sleeve garments, so we believe that SeamPose can be generalized to apparel with more seams, which offers more information on upper-body pose for the model to learn.

7.4 Different Seam Patterns in Everyday Apparel

As the ablation study shows, the performance decreased as we removed the sleeve electrode but the model can still infer the upper body pose from the seam electrodes around the shoulders. This indicates our approach holds the potential to work on short-sleeve shirts and sleeveless shirts with a slightly worse performance. However, if the shoulder electrodes do not exist, such as on strapped and strapless tops, it is unclear how our system will work. In the future, we plan to integrate sensing seams into more clothing patterns, including tops and bottoms (*i.e.*, pants, shorts, skirts). Further, we chose the zigzag stitch as a common stitch for seaming on home sewing machines. Zigzag stitches are stretchable and allow natural fabric stretches as the wearer moves. Stretching the fabric with sewn thread in any direction decreases the measured capacitances.

After hundreds of sessions with our prototype, there is no performance degradation caused by permanent deformations. From our experience, the stitch type may impact the raw signal qualities but is unlikely to greatly impact the tracking performance. Nevertheless, a formal characterization of stitch types (*e.g.*, overlock stitches common for industrial seaming) remains essential for future work.

8 CONCLUSION

We presented SeamPose, repurposing seams as capacitive sensors in a shirt for upper-body pose tracking. Without modifying the clothing surface, SeamPose integrates motion-capturing capabilities into clothing unobtrusively and "invisibly". In our proof-of-concept prototype, we machine-sewed conductive thread over existing seams in a long-sleeve shirt. The capacitive readings of the seam electrodes change as the wearer's body pose changes. Then the capacitive signals, acquired by an untethered sensing board, are processed by a customized deep-learning pipeline and continuously estimate 8 upper-body joint positions in 3D, relative to the pelvis. We evaluated SeamPose with a 12-participant user study and achieved a mean per joint position error (MPJPE) of 6.0 cm, comparable with that of related works, paving the way for everyday pose-tracking smart clothing.

ACKNOWLEDGMENTS

This project was partially supported by the National Science Foundation Grant Award No. 2239569. We thank Owen Koonce for appearing in the video figure. We also would like to thank the study participants and the reviewers.

REFERENCES

- [1] 2024. Meta Oculus. <https://www.meta.com/quest/>.
- [2] Talha Agcayazi, Jordan Tabor, Michael McKnight, Isaac Martin, Tushar K Ghosh, and Alper Bozkurt. 2020. Fully-textile seam-line sensors for facile textile integration and tunable multi-modal sensing of pressure, humidity, and wetness. *Advanced materials technologies* 5, 8 (2020), 2000155.
- [3] Karan Ahuja, Vivian Shen, Cathy Mengying Fang, Nathan Riopelle, Andy Kong, and Chris Harrison. 2022. Controllerpose: inside-out body capture with VR controller cameras. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [4] Roland Aigner, Andreas Pointner, Thomas Preindl, Rainer Danner, and Michael Haller. 2021. Txyz: Embroidering enameled wires for three degree-of-freedom mutual capacitive sensing. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [5] Riku Arakawa, Bing Zhou, Gurunandan Krishnan, Mayank Goel, and Shree K Nayar. 2023. MI-Poser: Human Body Pose Tracking Using Magnetic and Inertial Sensor Fusion with Metal Interference Mitigation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–24.
- [6] Asli Atalay, Vanessa Sanchez, Ozgur Atalay, Daniel M Vogt, Florian Haufe, Robert J Wood, and Conor J Walsh. 2017. Batch fabrication of customizable silicone-textile composite capacitive strain sensors for human motion tracking. *Advanced Materials Technologies* 2, 9 (2017), 1700136.
- [7] Hymalai Bello, Bo Zhou, Sungho Suh, and Paul Lukowicz. 2021. Mocapaci: Posture and gesture detection in loose garments using textile cables as capacitive antennas. In *Proceedings of the 2021 ACM International Symposium on Wearable Computers*. 78–83.
- [8] Sizhen Bian, Mengxi Liu, Bo Zhou, Paul Lukowicz, and Michele Magno. 2024. Body-Area Capacitive or Electric Field Sensing for Human Activity Recognition and Human-Computer Interaction: A Comprehensive Survey. *arXiv preprint arXiv:2401.06000* (2024).
- [9] Dario Cazzola, Timothy P Holsgrove, Ezio Preatoni, Harinderjit S Gill, and Grant Trewartha. 2017. Cervical spine injuries: a whole-body musculoskeletal model for the analysis of spinal loading. *PLoS one* 12, 1 (2017), e0169329.
- [10] Wenqiang Chen, Yexin Hu, Wei Song, Yingcheng Liu, Antonio Torralba, and Wojciech Matusik. 2024. CAvatar: Real-time Human Activity Mesh Reconstruction via Tactile Carpets. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 4 (2024), 1–24.

- [11] Xiaowei Chen, Xiao Jiang, Lishuang Zhan, Shihui Guo, Qunsheng Ruan, Guoliang Luo, Minghong Liao, and Yipeng Qin. 2023. Full-body human motion reconstruction with sparse joint tracking using flexible sensors. *ACM Transactions on Multimedia Computing, Communications and Applications* 20, 2 (2023), 1–19.
- [12] Jingyuan Cheng, Oliver Amft, and Paul Lukowicz. 2010. Active capacitive sensing: Exploring a new wearable sensing modality for activity recognition. In *International conference on pervasive computing*. Springer, 319–336.
- [13] Shreyosi Endow, Mohammad Abu Nasir Rakib, Anvay Srivastava, Sara Rastegarpouyani, and Cesar Torres. 2022. Embr: A Creative Framework for Hand Embroidered Liquid Crystal Textile Displays. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [14] Mohammad Iman Mokhelepour Esfahani and Maury A Nussbaum. 2018. A “smart” undershirt for tracking upper body motions: Task classification and angle estimation. *IEEE sensors Journal* 18, 18 (2018), 7650–7658.
- [15] Guido Gioberto, James Coughlin, Kaila Bibeau, and Lucy E Dunne. 2013. Detecting bends and fabric folds using stitched sensors. In *Proceedings of the 2013 International Symposium on Wearable Computers*. 53–56.
- [16] Shubham Goel, Georgios Pavlakos, Jathushan Rajasegaran, Angjoo Kanazawa, and Jitendra Malik. 2023. Humans in 4d: Reconstructing and tracking humans with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14783–14794.
- [17] Adam Goertz. 2020. A Capacitive Sensing Gym Mat for Exercise Classification & Tracking. (2020).
- [18] Jun Gong, Yu Wu, Lei Yan, Teddy Seyed, and Xing-Dong Yang. 2019. Tessuto: Contextual interactions on interactive fabrics with inductive sensing. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 29–41.
- [19] Chandramohan Gopalsamy, Sungmee Park, Rangaswamy Rajamanickam, and Sundaresan Jayaraman. 1999. The Wearable Motherboard™: The first generation of adaptive and responsive textile structures (ARTS) for medical applications. *Virtual Reality* 4 (1999), 152–168.
- [20] Tobias Grosse-Puppenthal, Christian Holz, Gabe Cohn, Raphael Wimmer, Oskar Bechtold, Steve Hodges, Matthew S Reynolds, and Joshua R Smith. 2017. Finding common ground: A survey of capacitive sensing in human-computer interaction. In *Proceedings of the 2017 CHI conference on human factors in computing systems*. 3293–3315.
- [21] Nur Al-huda Hamdan, Simon Voelker, and Jan Borchers. 2018. Sketch&stitch: Interactive embroidery for e-textiles. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [22] Yinghao Huang, Manuel Kaufmann, Emre Aksan, Michael J Black, Otmar Hilliges, and Gerard Pons-Moll. 2018. Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–15.
- [23] Jiayi Jiang, Paul Strel, Huajian Qiu, Andreas Fender, Larissa Laich, Patrick Snape, and Christian Holz. 2022. Avatarposer: Articulated full-body pose tracking from sparse motion sensing. In *European conference on computer vision*. Springer, 443–460.
- [24] Thorsten Karrer, Moritz Wittenhagen, Leonhard Lichtschlag, Florian Heller, and Jan Borchers. 2011. Pinstripe: eyes-free continuous input on interactive clothing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1313–1322.
- [25] Manuel Kaufmann, Yi Zhao, Chengcheng Tang, Lingling Tao, Christopher Twigg, Jie Song, Robert Wang, and Otmar Hilliges. 2021. Em-pose: 3d human pose estimation from sparse electromagnetic trackers. In *Proceedings of the IEEE/CVF international conference on computer vision*. 11510–11520.
- [26] Kushal Kedia, Prithwish Dan, Atiksh Bhardwaj, and Sanjiban Choudhury. 2023. Manicast: Collaborative manipulation with cost-aware human forecasting. *arXiv preprint arXiv:2310.13258* (2023).
- [27] Theodore T Kim, Mohamed A Zohdy, and Michael P Barker. 2020. Applying pose estimation to predict amateur golf swing performance using edge processing. *IEEE Access* 8 (2020), 143769–143776.
- [28] Ralf Kittler, Manfred Kayser, and Mark Stoneking. 2003. Molecular evolution of Pediculus humanus and the origin of clothing. *Current Biology* 13, 16 (2003), 1414–1417.
- [29] Chi-Jung Lee, Ruidong Zhang, Devansh Agarwal, Tianhong Catherine Yu, Vipin Gunda, Oliver Lopez, James Kim, Sicheng Yin, Boao Dong, Ke Li, Mose Sakashita, Francois Guimbretiere, and Cheng Zhang. 2024. EchoWrist: Continuous Hand Pose Tracking and Hand-Object Interaction Recognition Using Low-Power Active Acoustic Sensing On a Wristband. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 403, 21 pages. <https://doi.org/10.1145/3613904.3642910>
- [30] Ke Li, Ruidong Zhang, Siyuan Chen, Boao Chen, Mose Sakashita, François Guimbretière, and Cheng Zhang. 2024. EyeEcho: Continuous and Low-power Facial Expression Tracking on Glasses. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–24.
- [31] Ke Li, Ruidong Zhang, Bo Liang, François Guimbretière, and Cheng Zhang. 2022. EarIO: A Low-Power Acoustic Sensing Earable for Continuously Tracking Detailed Facial Movements. 6, 2, Article 62 (jul 2022), 24 pages. <https://doi.org/10.1145/3534621>
- [32] An Liang, Rebecca Stewart, Rachel Freire, and Nick Bryan-Kinns. 2021. Knit stretch sensor placement for body movement sensing. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. 1–7.
- [33] Hyunchul Lim, Yaxuan Li, Matthew Dressa, Fang Hu, Jae Hoon Kim, Ruidong Zhang, and Cheng Zhang. 2022. BodyTrak: Inferring Full-Body Poses from Body Silhouettes Using a Miniature Camera on a Wristband. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–21.
- [34] Ziyu Liu, Hongwen Zhang, Zhenghao Chen, Zhiyong Wang, and Wanli Ouyang. 2020. Disentangling and unifying graph convolutions for skeleton-based action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 143–152.
- [35] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics (TOG)* 34, 6 (2015), 1–16.
- [36] Yiyue Luo, Yunzhu Li, Pratyusha Sharma, Wan Shou, Kui Wu, Michael Foshey, Beichen Li, Tomás Palacios, Antonio Torralba, and Wojciech Matusik. 2021. Learning human–environment interactions using conformal tactile textiles. *Nature Electronics* 4, 3 (2021), 193–201.
- [37] Saif Mahmud, Ke Li, Guilin Hu, Hao Chen, Richard Jin, Ruidong Zhang, François Guimbretière, and Cheng Zhang. 2023. PoseSonic: 3D Upper Body Pose Estimation Through Egocentric Acoustic Sensing on Smartglasses. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–28.
- [38] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. 2015. Pose estimation for augmented reality: a hands-on survey. *IEEE transactions on visualization and computer graphics* 22, 12 (2015), 2633–2651.
- [39] Vimal Mollyn, Riku Arakawa, Mayank Goel, Chris Harrison, and Karan Ahuja. 2023. Imposer: Full-body pose estimation using imus in phones, watches, and earbuds. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [40] Hossein Mousavi Hondori and Maryam Khademi. 2014. A review on technical and clinical impact of microsoft kinect on physical therapy and rehabilitation. *Journal of medical engineering* 2014 (2014).
- [41] Movella. 2024. Xsens Products. <https://www.movella.com/products/xsens>.
- [42] Sara Nabil, Jan Kučera, Nikoleta Karastathi, David S Kirk, and Peter Wright. 2019. Seamless seams: Crafting techniques for embedding fabrics with interactive actuation. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 987–999.
- [43] OptiTrack. 2024. OptiTrack - Motion Capture Systems. <https://optitrack.com/>.
- [44] E Rehmi Post and Maggie Orth. 1997. Smart fabric, or “wearable clothing”. In *Digest of Papers. First International Symposium on Wearable Computers*. IEEE, 167–168.
- [45] Ivan Poupyrev, Nan-Wei Gong, Shiho Fukuhara, Mustafa Emre Karagozler, Carsten Schwesig, and Karen E Robinson. 2016. Project Jacquard: interactive digital textiles at scale. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 4216–4227.
- [46] Lala Shakti Swarup Ray, Bo Zhou, Sungho Suh, and Paul Lukowicz. 2023. Selecting the Motion Ground Truth for Loose-fitting Wearables: Benchmarking Optical MoCap Methods. In *Proceedings of the 2023 ACM International Symposium on Wearable Computers*. 27–32.
- [47] Olivia Ruston, Leon Watts, and Mike Fraser. 2021. More than it Seams: Garment Stitching in Wearable e-Textiles. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*. 1171–1182.
- [48] Christian Schöffmann, Zackory Erickson, and Hubert Zangl. 2022. CapSense: A real-time capacitive sensor simulation framework for physical human-robot interaction. *IEEE Robotics and Automation Letters* 7, 4 (2022), 9929–9936.
- [49] Qing Shuai, Chen Geng, Qi Fang, Sida Peng, Wenhao Shen, Xiaowei Zhou, and Hujun Bao. 2022. Novel view synthesis of human interactions from sparse multi-view videos. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–10.
- [50] Joshua R. Smith. 1996. Field mice: Extracting hand geometry from electric field measurements. *IBM systems journal* 35, 3.4 (1996), 587–608.
- [51] Jessica Stanley, John A Hunt, Phil Kunovski, and Yang Wei. 2022. A review of connectors and joining technologies for electronic textiles. *Engineering Reports* 4, 6 (2022), e12491.
- [52] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. 2019. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7728–7738.
- [53] Vicon. 2024. Vicon | Award Winning Motion Capture Systems. <https://www.vicon.com/>.
- [54] Jian Wang, Diogo Luvizon, Weipeng Xu, Lingjie Liu, Kripasindhu Sarkar, and Christian Theobalt. 2023. Scene-aware egocentric 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13031–13040.
- [55] Kai Wang, Xiaoyu Xu, Yiping Zheng, Da Zhou, Shihui Guo, Yipeng Qin, and Xiaohu Guo. 2023. Computational design of wiring layout on tight suits with

- minimal motion resistance. In *SIGGRAPH Asia 2023 Conference Papers*. 1–12.
- [56] Mark Weiser. 1991. The Computer for the 21 st Century. *Scientific american* 265, 3 (1991), 94–105.
- [57] Te-Yen Wu, Lu Tan, Yuji Zhang, Teddy Seyed, and Xing-Dong Yang. 2020. Capacitivo: Contact-based object recognition on interactive fabrics using capacitive sensing. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. 649–661.
- [58] Xinyu Yi, Yuxiao Zhou, Marc Habermann, Soshi Shimada, Vladislav Golyanik, Christian Theobalt, and Feng Xu. 2022. Physical inertial poser (pip): Physics-aware real-time human motion tracking from sparse inertial sensors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 13167–13178.
- [59] Tianhong Catherine Yu, Riku Arakawa, James McCann, and Mayank Goel. 2023. uKnit: A Position-Aware Reconfigurable Machine-Knitted Wearable for Gestural Interaction and Passive Sensing using Electrical Impedance Tomography. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [60] Tianhong Catherine Yu, Guilin Hu, Ruidong Zhang, Hyunchul Lim, Saif Mahmud, Chi-Jung Lee, Ke Li, Devansh Agarwal, Shuyang Nie, Jinseok Oh, et al. 2024. Ring-a-Pose: A Ring for Continuous Hand Pose Tracking. *arXiv preprint arXiv:2404.12980* (2024).
- [61] Tianhong Catherine Yu, Nancy Wang, Sarah Ellenbogen, and Cindy Hsin-Liu Kao. 2023. Skinergy: Machine-Embroidered Silicone-Textile Composites as On-Skin Self-Powered Input Sensors. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–15.
- [62] Dongzhi Zhang, Dongyue Wang, Zhenyuan Xu, Xixi Zhang, Yan Yang, Jingyu Guo, Bao Zhang, and Wenhao Zhao. 2021. Diversiform sensors and sensing systems driven by triboelectric and piezoelectric nanogenerators. *Coordination Chemistry Reviews* 427 (2021), 213597.
- [63] Bo Zhou, Daniel Geissler, Marc Faulhaber, Clara Elisabeth Gleiss, Esther Friederike Zahn, Lala Shakti Swarup Ray, David Gamarra, Vitor Fortes Rey, Sungho Suh, Sizhen Bian, et al. 2023. Mocapose: Motion capturing with textile-integrated capacitive sensors in loose-fitting smart garments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 1 (2023), 1–40.
- [64] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. 2019. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5745–5753.
- [65] Chengxu Zuo, Fang Jiawei, Shihui Guo, and Yipeng Qin. 2024. Self-adaptive motion tracking against on-body displacement of flexible sensors. *Advances in Neural Information Processing Systems* 36 (2024).
- [66] Chengxu Zuo, Yiming Wang, Lishuang Zhan, Shihui Guo, Xinyu Yi, Feng Xu, and Yipeng Qin. 2024. Loose inertial poser: Motion capture with IMU-attached loose-wear jacket. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

A USER STUDY VIDEO INSTRUCTION DETAILS

There are 54 instruction videos:

- 1-20: 20 gestures from [7] - lean forward, lean backward, lean to left, lean to right, turn left, turn right, shrug, pinch waist, forearm block, open arms, hands on the head, arms up, flappy bird, claps, walk ([7] indicates fake walk but we allowed the participant to locomote), butterfly swing, respect gesture, confuse gesture, frame picture, and stop gesture.
- 21: golf swing
- 22: right tennis swing
- 23: left tennis swing
- 24: right basketball dribble
- 25: left basketball dribble
- 26: basketball shooting
- 27: punches with alternate hands
- 28: left arm swing on the side of the body (like walking)
- 29: right arm swing on the side of the body (like walking)
- 30: left arm swing in front of the body
- 31: right arm swing in front of the body
- 32: head sequence 1 - move the head downward and upward

- 33: head sequence 2 - rotate the head clockwise and counter-clockwise
- 34: head sequence 2 - rotate the left and right
- 35: head sequence 3 - tilt the head from shoulder to shoulder
- 36: shoulder sequence 1 - move the shoulders up and down
- 37: shoulder sequence 2 - move the shoulders forward and backward
- 38: shoulder sequence 3 - rotate the shoulders forward and backward
- 39-41: arm sequence - with the arms down, the left/right/both arm(s) curl(s) from the inside, neutral, and outside tracks sequentially
- 42-44: arm sequence - with the arms open to the sides and parallel to the ground, the left/right/both arm(s) curl(s) from the inside, neutral, and outside tracks sequentially
- 45-47: arm sequence - with the arms front and parallel to the ground, the left/right/both arm(s) curl(s) from the inside, neutral, and outside tracks sequentially
- 48-50: arm sequence - with the arms raised over the head, left/right/both arm(s) curl(s) from the inside and neutral tracks sequentially (curing arms from the outside track are difficult to perform)
- 51-52: arm sequence - with the left/right arm open to the side, the right/left arm raises over the head and curl(s) from the inside and neutral tracks sequentially
- 53-54: arm sequence - with the left/right arm front and parallel to the ground, the right/left arm raises over the head and curl(s) from the inside and neutral tracks sequentially

Table 2: Anthropometric data of participants.

| | Arm Length (cm) | Bust (cm) | Waist (cm) | Height (cm) | Weight (kg) |
|------|--------------------|--------------|---------------|----------------|----------------|
| mean | 56.6 | 89.8 | 76.8 | 169.7 | 63.3 |
| std | 3.3 | 6.3 | 7.7 | 10.6 | 11.7 |
| max | 65.0 | 99.0 | 92.0 | 190.5 | 83.9 |
| min | 53.0 | 78.0 | 66.0 | 153 | 49.9 |