

EITWatch: Smartwatch-Integrated Planar Electrical Impedance Tomography for Hand Gesture Recognition

Xuanyou Liu*

Computer Science
Northwestern University
Evanston, Illinois, USA

xuanyouliu2030@u.northwestern.edu

Novel Alam*

Computer Science
Northwestern University
Evanston, Illinois, USA

novelalam2027@u.northwestern.edu

Karan Ahuja

Computer Science
Northwestern University
Evanston, Illinois, USA

kahuja@northwestern.edu

Abstract

Wrist Electrical Impedance Tomography (EIT) senses hand gestures from muscle- and tendon-driven impedance changes, but prior wrist-EIT systems require electrode coverage beyond the watch-back contact patch and separate analog front ends. We present EITWatch, the first wrist-EIT system built around smartwatch case-back geometry, asking whether this contact patch alone can support gesture recognition: eight planar electrodes in a 31 mm ring acquire 35 impedance measurements at 48 Hz. Because a planar array cannot encircle the wrist, EITWatch uses *multi-depth scanning* to sample multiple source-sink distances and current paths; it beat matched adjacent injection by 15.1/10.4 percentage points (macro/micro) across all 12 participants. In a prompted study, within-session leave-one-round-out accuracy reached 91.4%/92.5% (window/trial) for six macro-gestures, and 90.1%/91.5% (window/segment) for five micro-gestures plus *relax*; window-level cross-session and leave-one-user-out transfer reached 73.2%/70.4% and 63.1%/55.3% (macro/micro).

CCS Concepts

• **Human-centered computing** → **Gestural input**; *Ubiquitous and mobile computing*.

Keywords

surface-based planar EIT, smartwatch, micro-gesture recognition, wrist-worn sensing, wearable gesture input

ACM Reference Format:

Xuanyou Liu, Novel Alam, and Karan Ahuja. 2026. EITWatch: Smartwatch-Integrated Planar Electrical Impedance Tomography for Hand Gesture Recognition. In *The 39th Annual ACM Symposium on User Interface Software and Technology (UIST '26)*, November 02–05, 2026, Detroit, MI, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3830398.3830673>

1 Introduction

Smartwatches sit on the wrist all day, maintaining continuous skin contact through the case back, a contact that has rarely been used as an interactive surface. Electrical Impedance Tomography (EIT)

*Both authors contributed equally to this research. Xuanyou Liu is the corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *UIST '26, Detroit, MI, USA*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2856-3/26/11

<https://doi.org/10.1145/3830398.3830673>




Figure 1: EITWatch overview. (A) The custom PCB with eight planar EIT electrodes arranged in a 31 mm diameter ring on a 40 × 60 mm board. (B) The board mounted beneath a standard 40 mm watch case. (C) The system worn on the wrist during gesture input.

is well suited to exploit this contact: by injecting small alternating currents and measuring the resulting voltage patterns, it can detect impedance shifts caused by muscle and tendon deformation beneath the skin [10, 20]. Prior wrist-EIT systems show accurate gesture recognition, but they are worn as separate accessories, either circumferential bands [10, 20, 23] or a dedicated wristband carrying a localized array [24], and connected to external analog front ends. This leaves a direct HCI question unanswered: can the watch-back contact patch alone serve as a viable EIT sensing site for gesture recognition?

We present **EITWatch**, the first wrist-EIT system centered on smartwatch case-back geometry (Figure 1). The prototype features a custom PCB with eight skin-facing planar electrodes arranged in a 31 mm ring; an onboard ESP32-S3 microcontroller sequences these electrodes to acquire 35 impedance measurements at a frame rate of 48 Hz. The current discrete analog front end extends 20 mm beyond the case, so the contribution is not a fully miniaturized smartwatch; rather, it isolates the watch-back sensing geometry and tests whether that contact area alone provides discriminative impedance measurements.

Moving to a planar watch-back geometry also changes the acquisition design. A dorsal array cannot surround the wrist cross-section, so the adjacent-injection schedule used by circumferential systems mainly emphasizes shallow tissue. EITWatch therefore uses *multi-depth scanning*, a multi-distance, multi-path protocol

Table 1: Hardware and task scope of electrical impedance sensing systems for hand gesture recognition. Circ.: circumferential; Local 2×2: a localized 4-electrode array carried on a wristband rather than encircling the wrist. Frame rates are as reported by each system and are not directly comparable, since several scale with electrode count.

	Tomo [20]	HighRes [21]	EIT-kit [23]	EITPose [10]	EI-Lite [24]	BandEI [19]	EITWatch
Electrodes	8	8 / 16 / 32	Up to 64	8	4	32	8
Frame rate (Hz)	10	3–100	≤10	10	100	30	48
Electrode placement	Circ.	Circ.	Circ.	Circ.	Local 2×2	Finger + wrist	Planar (dorsal)
Form factor	Forearm band	Forearm band	Configurable	Wrist band	Wrist band	Flex. bandage	Watch back
Gesture task	Macro	Macro	Configurable	Macro	Micro	Macro + micro	Macro + micro

that sweeps source-sink separation and direction within each frame. In a homogeneous half-space, its source-sink distances correspond to heuristic characteristic depth scales from approximately 6 mm to 15.5 mm (Section 3); on the heterogeneous wrist, changing the sink also changes the angular path through muscles and tendons. We validate the protocol in a water tank and through a matched on-body ablation.

Across 12 participants, within-session leave-one-round-out evaluation exceeded 90% window-level accuracy for both gesture sets, measuring within-session repeatability rather than deployable recognition; 48-hour cross-session and leave-one-user-out transfer dropped substantially, and a matched ablation favored multi-depth scanning over adjacent injection on the same hardware and tasks (Section 6). Together, these results indicate that watch-back EIT carries discriminative signal, while re-donning and person-independent transfer remain open challenges.

The contributions of this paper are (1) **the EITWatch device concept**: the first wrist-EIT system built around a standard smartwatch case-back contact area, isolating the sensing geometry to the watch-back contact patch even though the current electronics are not yet fully miniaturized; and (2) **feasibility evidence**: a 12-participant evaluation of within-session, cross-session, and cross-user recognition, a matched comparison showing that multi-depth scanning outperforms adjacent injection on this planar system, and measurements of prototype latency and power.

To support future research in planar wrist-EIT, we release the EITWatch schematics, PCB layouts, and firmware at https://github.com/XuanyouLiu/EITWatch_Hardware.

2 Related Work

Wrist-worn EIT. Prior wrist-EIT systems show that the wrist exhibits rich impedance structure for gesture recognition, but all require hardware beyond the watch-back contact patch. Tomo [20], HighRes [21], and EIT-kit [23] use circumferential forearm or wrist bands; EITPose [10] and EI-Lite [24] reduce electrode count but still rely on a separate band. In each case, the electrode array connects to an external analog front end board that often remains larger than the wearable band alone. BandEI [19] explores flexible placement including the wrist, yet performs best with 24 finger channels rather than 8 wrist-only channels. EITWatch instead places all electrodes on the case back, asking whether a watch-back-sized planar array can support gesture recognition at all. Table 1 situates hardware and task scope only; the paper’s main evidence comes from matched experiments on the same device, especially the protocol ablation in Section 6.3.

Smartwatch gesture input. Non-EIT smartwatch input includes EMG [16], wrist gestures [6], bio-acoustics [11], ultrasonic beam-forming [8], active sonar [15], body-coupled RF [22], and depth sensing [18]. Most require additional hardware or active signal emission beyond the watch itself; EITWatch instead asks whether the existing case-back contact patch can become a sensing surface in its own right.

Surface planar EIT. Planar EIT places all electrodes on one surface and senses into the volume beneath it. Prior work established reconstruction algorithms [14], demonstrated probe-based planar sensing at depths on the order of 2–4 cm in controlled phantoms [9], and characterized depth-dependent resolution for one-sided geometries [4]. That literature makes watch-back sensing physically plausible, but does not address smartwatch-scale planar arrays for gesture recognition. EITWatch addresses that device-level question directly.

3 Planar Sensing on the Watch Back

EITWatch places all eight electrodes on the dorsal wrist surface within a 31 mm ring. To understand what this compact planar array can sense, we examine how current penetration varies with electrode spacing.

When a small alternating current is injected between two surface electrodes separated by distance L , the resulting electric potential ϕ inside the tissue domain Ω satisfies the quasi-static Laplace equation [3, 9]:

$$\nabla \cdot (\sigma \nabla \phi) = 0 \quad \text{in } \Omega \quad (1)$$

where σ is the spatial conductivity distribution. The current density $\mathbf{J} = -\sigma \nabla \phi$ does not flow in a straight line between the electrodes; instead, it curves downward into the tissue volume. For a homogeneous half-space, the current density magnitude along the depth axis z is:

$$|\mathbf{J}(z)| \propto \frac{1}{((L/2)^2 + z^2)^{3/2}} \quad (2)$$

This solution shows that current density decays continuously with depth. We use $L/2$ as a heuristic characteristic depth scale, rather than a discrete penetration boundary [7]. As shown in Figure 2, adjacent electrodes ($L \approx 11.9$ mm) weight shallower regions more strongly, whereas diametrically opposite electrodes ($L \approx 31.0$ mm) extend the field farther into the half-space. This homogeneous model isolates the effect of separation; it does not model wrist anatomy.

Gesture recognition exploits the impedance modulation produced as muscle contracts and its fibers shorten and thicken. Surface




Figure 2: Current penetration in a homogeneous half-space. Varying the injection-pair distance L from adjacent to opposite increases the heuristic characteristic depth scale ($L/2$); current density nevertheless decays continuously with depth.

measurements on the anterior forearm at 50 kHz, EITWatch’s operating frequency, report that resistance *increases* under isometric finger-flexor contraction [17], and localized impedance myography of the biceps places the magnitude near 10% at maximal voluntary contraction [12]. This modulation acts against a static background of poorly conducting tissue, including cortical bone (~ 0.02 S/m) and fat (~ 0.02 – 0.04 S/m) [5]. Because EITWatch senses from the dorsal surface only, bone and fat steer the field and create position-dependent sensitivity rather than the more uniform wrap-around view of circumferential systems. Multi-depth scanning varies both separation and sink direction, thereby sampling multiple depth profiles and angular paths within each frame. The water-tank study below compares the complete protocols under controlled conditions, while the on-body ablation evaluates them on real wrist anatomy; neither comparison isolates separation from angular coverage or channel count.

4 EITWatch System Design

4.1 Electrode Array

Eight gold-plated stainless steel disk electrodes (diameter 2 mm) are mounted on the skin-facing surface of a custom PCB (Figure 1A). They are evenly distributed on a $\varnothing 31$ mm circle, yielding ~ 11.9 mm adjacent spacing. Because the ring fits within the 40 mm watch case back, the board attaches directly to the case and presses all electrodes flush against the dorsal wrist.

4.2 PCB and Firmware

Figure 3 shows the circuit architecture. The custom six-layer PCB (40×60 mm) carries eight gold electrodes on the bottom and mixed-signal conditioning on top. The electrode ring and switch matrices occupy the 40 mm region beneath the watch case, while the remaining 20 mm extension houses the discrete analog front end. This partition matches the watch-back footprint while keeping the prototype debuggable with off-the-shelf components before full electronics miniaturization. An AD5930 [2] waveform generator produces a 50 kHz sinusoid that is amplified and converted to a current source by an ADA4841-based voltage-to-current stage. An AD5270 digital potentiometer sets the injected current to approximately 1 mA peak-to-peak. On the sensing side, an AD8220 instrumentation amplifier with AD5270-adjustable gain amplifies the




Figure 3: EITWatch circuit architecture. An ESP32-S3 microcontroller sequences eight electrodes through a 35-channel multi-depth scanning schedule at ~ 48 Hz.

differential voltage, which is digitized by a 12-bit ADC (AD7450 [1]). Four 8-channel matrix switches (ADG738) route injection and sense pairs to the electrodes for the multi-depth schedule (Section 4.3).

A Seeed XIAO ESP32-S3 microcontroller running FreeRTOS sequences the scanning schedule at 48 Hz, chosen to resolve the sub-second transients that characterize micro-gestures. For the gesture recognition study (Section 5), the device streams the 35 scalar features over Wi-Fi to a host computer for logging and offline classification.

4.3 Measurement Configuration

As a baseline we use four-pole adjacent injection, the standard in circumferential wrist EIT: the source-sink pair occupies neighboring electrodes and steps around the ring, with voltage sensed differentially on non-drive electrodes. With 8 electrodes and 5 non-drive differential measurements per drive pair, this yields a 40-channel baseline used in Section 6.3. On a dorsal planar array, however, this protocol is shallow, concentrating current near the skin surface ($z \approx L/2 \approx 6$ mm) and reducing sensitivity to deeper muscle deformation.

We therefore use a *multi-depth scanning* protocol. The electrode at the 12-o’clock position serves as the fixed positive source; the negative sink iterates through the remaining 7 positions, producing injection pairs from adjacent ($L \approx 11.9$ mm) to diametrically opposite ($L \approx 31.0$ mm). Five four-pole differential measurements per injection pair yield $7 \times 5 = 35$ measurements per frame, so only the injection geometry differs from the baseline. The schedule samples




Figure 4: Water-tank validation across four anomaly depths (2, 5, 10, and 20 mm). Multi-depth scanning retains stronger deep-anomaly contrast than adjacent injection at 10–20 mm, consistent with greater deep-anomaly sensitivity but not isolating depth as the cause.

multiple source-sink distances and angular paths in one cycle. A pilot comparison across all eight possible source electrodes on two participants showed less than 1.5 percentage points of variation in LORO accuracy, indicating that source choice does not meaningfully affect recognition on this symmetric ring layout. This schedule trades a small reduction in per-frame channel count for broader multi-depth, multi-path coverage.

4.4 Water-Tank Validation

We compared the complete protocols in a saline tank with a conductive anomaly at 2, 5, 10, and 20 mm from the electrode surface (Figure 4), using identical hardware and 10 repeats per depth. For each trial, we reconstructed the differential signal with pyEIT [13] and quantified it as the mean absolute reconstructed value within a fixed ROI centered on the anomaly location. For each depth, responses were averaged over repeats and normalized by the corresponding 2 mm value within the same protocol. Both protocols detected shallow anomalies similarly, with normalized response near 1.00 at 2 mm and about 0.80 at 5 mm. At greater depths, however, the multi-depth protocol retained a larger normalized response: at 10 mm, 0.58 ± 0.04 versus 0.37 ± 0.05 for adjacent injection, and at 20 mm, 0.31 ± 0.03 versus 0.14 ± 0.02 . This result is consistent with stronger sensitivity to deeper anomalies under multi-depth scanning. However, because the protocols also differ in source-sink separation, angular current paths, and per-frame channel count (35 vs. 40), it does not isolate depth as the cause of the difference.

5 Gesture Recognition Study

5.1 Gesture Set Design

We evaluated two gesture vocabularies (Figure 5).

Macro-Gestures. Six macro-gestures were used: **Six**, **Gun**, **Point**, **Thumb Up**, **Stretch**, and **Fist**. Three of these (Thumb Up, Stretch, and Fist) are drawn from the gesture set of Tomo [20].

Micro-Gestures. Five micro-gestures were selected as subtle, eyes-free smartwatch commands: **Swipe Left** and **Swipe Right** (thumb slides across the index finger), **Pinch** (two rapid thumb-to-index pinches), **Splay** (rapid full-hand opening), and **Wrist Flip** (rapid forearm rotation $\approx 90^\circ$). A *relax* class from all untagged intervals

yields a six-class task on prompted recordings, letting us evaluate both gesture discrimination and relax-state rejection without claiming open-world online spotting.

5.2 Participants and Apparatus

Twelve right-handed participants (6F, 6M; age 20–31; wrist circumference 181 ± 10.2 mm; BMI 25.3 ± 4.8 kg/m²) were recruited from the university community (IRB-approved; informed consent obtained). EITWatch was positioned on the distal dorsal wrist above the ulnar styloid and secured with a silicone strap without conductive gel; participants tightened it until the watch no longer slid during use. We did not perform additional contact-quality, pressure, or per-user firmware calibration so that performance would reflect the fixed watch-back geometry and acquisition protocol rather than user-specific compensation. A Meta Quest 3 provided VR prompts and optical hand-tracking verification.

5.3 Procedure

Each participant completed a ~ 30 min session in two blocks. Before recording, participants practiced with VR prompts and hand-tracking feedback. In the *macro-gesture block*, 10 rounds of 6 gestures were presented in randomized order; the Quest 3 verified the pose, each gesture was held for 2 s while EIT data were recorded, and a 10 s break separated rounds. Because every macro trial was a prompted held pose, this closed-set block contained no idle interval to label as *relax*. In the *micro-gesture block*, 10 rounds of the 5 prompted transient gestures were collected, with 5 consecutive repetitions per round. Participants self-tagged onset/offset via spacebar, and a 5 s break separated rounds. All untagged intervals, including inter-repetition and inter-round periods, were labeled *relax*, which forms the sixth class and constitutes 51.7% of micro-gesture windows. Separately, all participants completed a new study under the adjacent-injection protocol for the matched acquisition-protocol comparison. All participants returned 48 hours later for a second session after independently re-donning the device.

5.4 Feature Extraction and Classification

Raw 35-channel EIT frames were low-pass filtered with a zero-phase 4th-order Butterworth filter (0.5 Hz macro / 1.0 Hz micro cutoff) over each continuous recording before the LORO split, then IQR filtered ($k=3$) for outliers. The lower macro cutoff retains slowly varying held poses, whereas the higher micro cutoff retains faster transient motion. Because the inter-round breaks exceed the filters'




Figure 5: Gesture set. Top row: six static macro-gestures. Bottom row: six dynamic micro classes (including *relax*); arrows indicate motion direction.

Table 2: Per-participant recognition accuracy (%). Macro/Micro LORO use within-session multi-depth scanning; Trial and Segment are majority votes over a prompted hold or an annotated interval. Cross-session, LOUO, and Adjacent are *window-level*; Adjacent is matched LORO under 40-channel adjacent injection. Mean $\pm$: across-participant sample SD. Round SD: mean of per-participant fold SDs.

ID	Macro LORO (multi-depth)			Micro LORO (multi-depth)			Cross-session (window)		LOUO (window)		Adjacent (window)	
	Window	Trial	Round SD	Window	Segment	Round SD	Mac	Mic	Mac	Mic	Mac	Mic
P1	93.0	94.0	10.3	91.4	94.5	3.3	67.6	78.0	64.9	58.8	78.4	85.4
P2	91.2	93.0	8.6	89.6	91.1	6.1	83.1	59.9	69.3	47.6	72.0	82.9
P3	95.8	95.7	6.0	93.6	95.4	2.8	78.2	63.6	59.6	40.8	77.0	87.2
P4	88.9	90.6	12.6	88.6	89.2	8.2	72.6	71.5	59.9	68.7	76.2	79.8
P5	86.3	87.9	16.6	87.7	85.9	7.5	64.9	61.7	59.8	46.7	72.9	71.7
P6	92.5	93.6	9.0	90.2	92.9	5.5	74.5	77.9	71.2	62.8	79.9	84.9
P7	89.4	91.0	10.9	89.2	90.0	6.3	84.3	79.4	54.7	59.8	72.9	75.0
P8	89.6	91.4	12.2	89.3	91.0	6.8	71.3	66.5	48.4	50.1	75.3	73.9
P9	94.8	94.4	6.0	91.5	94.2	2.8	79.8	77.0	76.4	63.5	77.7	81.8
P10	92.0	93.1	12.4	90.2	92.0	5.7	69.9	76.8	62.5	54.8	84.0	81.7
P11	95.6	95.1	6.1	91.7	93.4	2.8	65.0	64.1	63.7	52.1	75.7	79.1
P12	87.7	90.2	14.1	88.2	88.4	9.4	67.2	68.4	66.8	57.9	73.6	73.0
Mean	91.4 $\pm$ 3.1	92.5 $\pm$ 2.3	10.4	90.1 $\pm$ 1.7	91.5 $\pm$ 2.8	5.6	73.2 $\pm$ 6.8	70.4 $\pm$ 7.2	63.1 $\pm$ 7.5	55.3 $\pm$ 8.1	76.3 $\pm$ 3.4	79.7 $\pm$ 5.2

effective span, this pre-split filtering cannot carry gesture information across a LORO boundary, and no windows overlap between training and held-out rounds. Macro-gestures used sliding windows of 40 frames with step size 2 (~ 0.83 s at 48 Hz); micro-gestures used 3 frames with step size 1 (~ 0.06 s). From each window we computed five statistics per channel (mean, SD, slope, range, and net displacement), giving 175 features per window (35×5). An Extra-Trees classifier (200 estimators macro, 400 micro; balanced class weighting) was used. IQR filtering, feature extraction, and model selection were performed independently within each fold; window, cutoff, and tree-count settings were fixed using a validation round from the training data before the test round was scored. LORO is therefore a user-dependent within-session repeatability metric, not a deployment estimate. We report accuracy at three granularities: *window* (each classifier decision), *trial* (majority vote over windows in a prompted macro hold), and *segment* (majority vote over windows in an annotated micro-gesture or contiguous relax interval). Window-level accuracy is the primary per-decision metric because trial and segment voting assume known interval boundaries. This study evaluates classification on prompted recordings rather than open-world onset detection.

6 Results

We evaluate EITWatch along three axes: within-session discriminability, robustness across sessions and users, and protocol choice, then report the on-device latency and power cost. Table 2 reports per-participant accuracies and foregrounds the distinction between user-dependent within-session results and the harder transfer settings.

6.1 Within-Session Gesture Recognition

For macro-gestures, mean LORO accuracy was 91.4% (SD = 3.1) per window and 92.5% (SD = 2.3) after trial voting, with a mean per-round SD of 10.4 across leave-one-round folds. For the prompted six-class micro task, it was 90.1% (SD = 1.7) per window and 91.5% (SD = 2.8) after segment-level voting within annotated segments,




Figure 6: Pooled leave-one-round-out confusion matrices (N = 12). Left: macro-gestures at the *trial* level. Right: micro-gestures at the *segment* level. Cells report counts and row-normalized percentages.

with a mean per-round SD of 5.6. Figure 6 shows the pooled macro trial-level and micro segment-level confusion matrices across all 12 participants (macro Acc = 92.5%, F1 = 92.4%; micro Acc = 91.5%, F1 = 91.4%). Stretch and Thumb Up were recognized perfectly; Point was the weakest macro class (84%), mainly confused with Fist or Thumb Up. On the micro side, Relax and Wrist Flip reached 99%, while Swipe Right was hardest (82%), mainly confused with Splay and Swipe Left. These within-session results measure repeatability within a single donning session; they do not establish person-independent or open-world performance. The micro *relax* class accounts for 51.7% of windows; balanced class weights and the primary window-level metric limit the influence of this imbalance.

6.2 Cross-Session and Cross-User Robustness

All 12 participants returned after 48 hours and re-donned EITWatch unassisted. For each participant, we trained a model on all first-session data and evaluated it on that participant's complete second-session recording without retraining. This yielded 73.2% (SD = 6.8) macro and 70.4% (SD = 7.2) micro window-level accuracy. Leave-one-user-out window-level evaluation yielded 63.1% (SD = 7.5) and 55.3% (SD = 8.1), respectively. These transfer results

remain above six-class chance but are substantially below within-session LORO. Re-donning changes contact location, pressure, and impedance, while anatomical variation changes the planar current paths; person-independent transfer is therefore the hardest evaluated regime. The gap bounds the current system's practical generalization rather than supporting a deployment-ready claim.

6.3 Protocol Ablation

All 12 participants completed a matched comparison under a standard 40-channel adjacent-injection protocol on the same device and pipeline, which reached window-level LORO accuracy of 76.3% (SD = 3.4) for macro- and 79.7% (SD = 5.2) for micro-gestures. Multi-depth scanning improved on this by 15.1 and 10.4 percentage points despite using fewer channels per frame (35 vs. 40), and every participant improved. This supports multi-depth, multi-path scanning as a stronger starting point for the planar watch back, though it does not establish optimality across planar EIT configurations.

6.4 Latency and Power

To approximate a more realistic on-watch deployment, we separately built an online dual-core pipeline that runs entirely on the ESP32-S3: one core handles acquisition and rolling-window updates, and the other runs preprocessing and ExtraTrees inference. Predictions update in lockstep with the rolling window. The latency and power figures below are measured on this on-device path, not on the Wi-Fi logging path used in the user study. Frames arrive every 20.8 ms at 48 Hz. The measure-to-inference update takes 3.07 ms for macro-gestures (1.01 ms preprocessing and 2.06 ms inference, SD = 0.05 ms) and 5.25 ms for micro-gestures (0.75 ms preprocessing and 4.50 ms inference, SD = 0.03 ms), both fitting within one rolling-window update interval. End-to-end motion-to-photon latency averages 847.1 ms for macro-gestures, dominated by the 0.83 s feature window, and 72.9 ms for micro-gestures. Running acquisition and inference continuously draws 35 mA at 4.3 V (0.15 W), corresponding to approximately 8.6 h from an idealized 300 mAh, 4.3 V battery. These are prototype measurements on the XIAO ESP32-S3, not an all-day smartwatch result.

7 Limitations and Discussion

The central finding is that the planar watch-back patch carries gesture-discriminative EIT signal: within-session performance held across all 12 participants, multi-depth scanning beat adjacent injection for every participant, and transfer across sessions and users stayed above chance. The much lower transfer accuracies matter equally: EITWatch is a feasibility prototype, not a deployment-ready recognizer.

Its limitations set the agenda. We do not sweep strap force, watch rotation, contact impedance, or wrist tissue composition, so robustness is only indirectly probed. The prompted micro task relies on annotated intervals and a small vocabulary, leaving open-world spotting, false-activation rates, and live UI feedback for separate study. And while inference fits each 20.8 ms frame, macro latency is dominated by its 0.83 s window.

These point to clear next steps: integrate the front end into a watch case, improve transfer via don-time calibration or user-adaptive training, and broaden gesture sets and populations.

References

- [1] Analog Devices Inc. 2009. AD7450: Differential Input, 1 MSPS 12-Bit ADC in μ SOIC-8 and SO-8. Product datasheet. <https://www.analog.com/en/products/ad7450.html>
- [2] Analog Devices Inc. 2013. AD5930: Programmable Frequency Sweep and Output Burst Waveform Generator. Product datasheet. <https://www.analog.com/en/products/ad5930.html>
- [3] Richard H. Bayford. 2006. Bioimpedance Tomography (Electrical Impedance Tomography). *Annual Review of Biomedical Engineering* 8, 1 (2006), 63–91. doi:10.1146/annurev.bioeng.8.061505.095716
- [4] Bo Chen and Manuchehr Soleimani. 2019. Depth Analysis of Planar Array for 3D Electrical Impedance Tomography. *IEEE Sensors Journal* 19, 22 (2019), 10710–10718. doi:10.1109/JSEN.2019.2929625
- [5] S. Gabriel, R. W. Lau, and C. Gabriel. 1996. The dielectric properties of biological tissues: III. Parametric models for the dielectric spectrum of tissues. *Physics in Medicine and Biology* 41, 11 (1996), 2271–2293. doi:10.1088/0031-9155/41/11/003
- [6] Jun Gong, Xing-Dong Yang, and Pourang Irani. 2016. WristWhirl: One-handed Continuous Smartwatch Input using Wrist Gestures. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 861–872. doi:10.1145/2984511.2984563
- [7] Sverre Grimnes and Ørjan G. Martinsen. 2008. *Bioimpedance and Bioelectricity Basics* (2nd ed.). Academic Press, London. doi:10.1016/b978-0-12-374004-5.x0001-3
- [8] Yasha Iravantchi, Mayank Goel, and Chris Harrison. 2019. BeamBand: Hand Gesture Sensing with Ultrasonic Beamforming. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA, 1–10. doi:10.1145/3290605.3300245
- [9] Tzu-Jen Kao, David Isaacson, Jonathan C. Newell, and Gary J. Saulnier. 2006. A 3D reconstruction algorithm for EIT using a handheld probe for breast cancer detection. *Physiological Measurement* 27, 5 (2006), S1–S11. doi:10.1088/0967-3334/27/5/S01
- [10] Alexander Kyu, Hongyu Mao, Junyi Zhu, Mayank Goel, and Karan Ahuja. 2024. EITPose: Wearable and Practical Electrical Impedance Tomography for Continuous Hand Pose Estimation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. ACM, New York, NY, USA, 1–10. doi:10.1145/3613904.3642663
- [11] Gierad Laput, Robert Xiao, and Chris Harrison. 2016. ViBand: High-Fidelity Bio-Acoustic Sensing Using Commodity Smartwatch Accelerometers. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 321–333. doi:10.1145/2984511.2984582
- [12] Le Li, Henry Shin, Xiaoyan Li, Sheng Li, and Ping Zhou. 2016. Localized Electrical Impedance Myography of the Biceps Brachii Muscle during Different Levels of Isometric Contraction and Fatigue. *Sensors* 16, 4 (2016), 581. doi:10.3390/s16040581
- [13] Benyuan Liu, Bin Yang, Canhua Xu, Junying Xia, Meng Dai, Zhenyu Ji, Fusheng You, Xiuzhen Dong, Xuetao Shi, and Feng Fu. 2018. pyEIT: A Python Based Framework for Electrical Impedance Tomography. *SoftwareX* 7 (2018), 304–308. doi:10.1016/j.softx.2018.09.005
- [14] Jennifer L. Mueller, David Isaacson, and Jonathan C. Newell. 1999. A reconstruction algorithm for electrical impedance tomography data collected on rectangular electrode arrays. *IEEE Transactions on Biomedical Engineering* 46, 11 (1999), 1379–1386. doi:10.1109/10.797998
- [15] Rajalakshmi Nandakumar, Vikram Iyer, Desney Tan, and Shyamnath Gollakota. 2016. FingerIO: Using Active Sonar for Fine-Grained Finger Tracking. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1515–1525. doi:10.1145/2858036.2858580
- [16] T. Scott Saponas, Desney S. Tan, Dan Morris, Ravin Balakrishnan, Jim Turner, and James A. Landay. 2009. Enabling Always-Available Input with Muscle-Computer Interfaces. In *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology (UIST '09)*. ACM, New York, NY, USA, 167–176. doi:10.1145/1622176.1622208
- [17] Carl A. Shiffman, Ronald Aaron, and Seward B. Rutkove. 2003. Electrical impedance of muscle during isometric contraction. *Physiological Measurement* 24, 1 (2003), 213–234. doi:10.1088/0967-3334/24/1/316
- [18] Srinath Sridhar, Anders Markussen, Antti Oulasvirta, Christian Theobalt, and Sebastian Boring. 2017. WatchSense: On- and Above-Skin Input Sensing through a Wearable Depth Sensor. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3891–3902. doi:10.1145/3025453.3026005
- [19] Hongrui Wu, Feier Long, Hongyu Mao, JaeYoung Moon, Junyi Zhu, and Yiyue Luo. 2025. BandEI: A Flexible Electrical Impedance Sensing Bandage for Deep Muscles and Tendons. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. ACM, New York, NY, USA, 1–12. doi:10.1145/3746059.3747732
- [20] Yang Zhang and Chris Harrison. 2015. Tomo: Wearable, Low-Cost Electrical Impedance Tomography for Hand Gesture Recognition. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software and Technology (UIST '15)*.

- ACM, New York, NY, USA, 167–173. doi:10.1145/2807442.2807480
- [21] Yang Zhang, Robert Xiao, and Chris Harrison. 2016. Advancing Hand Gesture Recognition with High Resolution Electrical Impedance Tomography. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 843–850. doi:10.1145/2984511.2984574
- [22] Yang Zhang, Junhan Zhou, Gierad Laput, and Chris Harrison. 2016. SkinTrack: Using the Body as an Electrical Waveguide for Continuous Finger Tracking on the Skin. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1491–1503. doi:10.1145/2858036.2858082
- [23] Junyi Zhu, Jackson C. Snowden, Joshua Verdejo, Emily Chen, Paul Zhang, Hamid Ghaednia, Joseph H. Schwab, and Stefanie Mueller. 2021. EIT-kit: An Electrical Impedance Tomography Toolkit for Health and Motion Sensing. In *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology (UIST '21)*. ACM, New York, NY, USA, 400–413. doi:10.1145/3472749.3474758
- [24] Junyi Zhu, Tianyu Xu, Jiayu Wang, Emily Guan, JaeYoung Moon, Stiven Morvan, D. Shin, Andrea Colaço, Stefanie Mueller, Karan Ahuja, Yiyue Luo, and Ishan Chatterjee. 2025. EI-Lite: Electrical Impedance Sensing for Micro-gesture Recognition and Pinch Force Estimation. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. ACM, New York, NY, USA, 1–14. doi:10.1145/3746059.3747671