TacPose: Continuous Hand Pose Estimation from Wrist-Worn Triaxial Tactile Sensing

Tai Inui
Waseda University
Tokyo, Japan
taiinui556@suou.waseda.jp

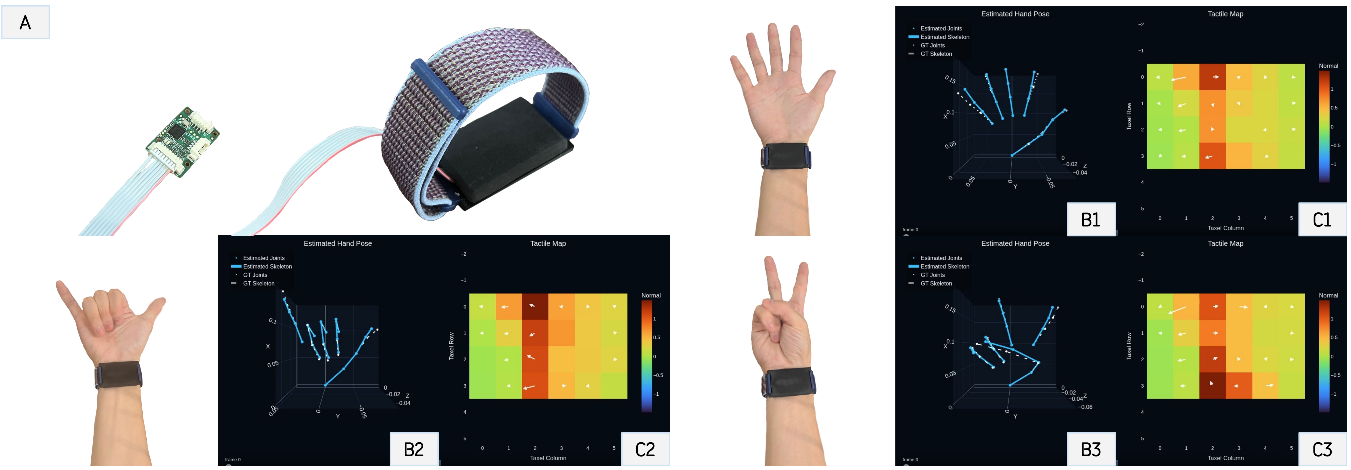

Figure 1: TacPose overview. TacPose is a soft, wrist-worn sensing device for continuous hand pose estimation (A). The system uses a Hall-effect-based tactile sensor mounted on the inner wrist to capture subtle skin, tendon, and tissue deformations caused by hand and finger motion. These deformations are represented as dense triaxial tactile force matrices (C), which are then mapped to continuous 3D hand pose estimates (B).

Abstract

Hands are central to embodied interaction, robot teleoperation, and learning from human demonstration. Existing hand pose tracking systems are often vision-based, glove-based, or electrode-based, which can introduce limitations related to occlusion, privacy, comfort, or setup burden. This paper presents TacPose, a wrist-worn tactile sensing approach for continuous hand pose estimation. TacPose uses a compact 4×6 triaxial tactile sensor array mounted on the inner wrist to measure local skin, tendon, and soft-tissue deformation caused by hand and finger motion. A spatial-temporal learning pipeline maps short windows of tactile measurements to wrist-relative 3D hand landmarks. In a preliminary study with five participants, TacPose achieved a within-session mean per-joint position error (MPJPE) of 3.44 mm. Cross-session and cross-user evaluation increased error to 12.27 mm and 13.76 mm, respectively, indicating that sensor placement, strap tightness, and pressure baseline are major generalization challenges. Ablation studies show that normal-force-only input generalizes better than full triaxial input in cross-user evaluation, while short temporal windows and a spatial temporal convolutional architecture improve performance. These results suggest that wrist-worn tactile sensing is a promising complementary modality for low-obtrusion, privacy-preserving hand pose estimation in HCI, XR, rehabilitation, and robot teleoperation.

CCS Concepts

• **Human-centered computing → Interaction devices**.

Keywords

Tactile Sensing, Hand Pose, Hand Gesture, Input, Interaction Technique, Extended Reality

1 Introduction

Human hands are expressive, dexterous, and central to how people manipulate objects, communicate intent, and interact with digital systems. Continuous hand pose estimation enables virtual object manipulation, mid-air interaction, rehabilitation monitoring, sign-language-related input, robot teleoperation, and learning from human demonstration. Unlike discrete gesture recognition, continuous pose estimation preserves the geometry of the hand, making it useful for interaction and robotics tasks that depend on gradual finger motion, grasp preshaping, and dexterous manipulation [1, 4].

Vision-based hand tracking has enabled substantial progress and is now common in XR systems and real-time tracking pipelines [20]. However, cameras remain sensitive to self-occlusion, field-of-view limitations, lighting, object occlusion, and social privacy concerns. Wearable approaches such as EMG, RF, acoustic sensing, wrist cameras, and pressure sensing address some of these issues, but introduce their own tradeoffs in setup burden, signal stability, hardware complexity, or wearability [6, 9, 11, 13, 16].

This paper explores a mechanically direct alternative: estimating continuous hand pose from tactile deformation at the wrist. TacPose uses a commercially available 4×6 XELA uSkin triaxial tactile sensor mounted to the inner wrist. The sensor measures distributed normal and tangential deformation caused by tendon glide, skin stretch, and soft-tissue displacement during hand motion. A supervised model maps short tactile windows to wrist-relative 3D hand landmarks.

The central research question is practical and empirical: can a compact wrist-mounted tactile array recover continuous hand pose, and what design factors affect generalization across sessions and users? Our preliminary evaluation shows that TacPose can estimate hand pose accurately when the wearing condition is consistent, but also reveals that placement and pressure variation strongly affect robustness.

This paper contributes:

- A compact wrist-worn triaxial tactile sensing approach for continuous 3D hand pose estimation.
- A spatial-temporal learning pipeline that preserves tactile grid structure and predicts wrist-relative hand landmarks.
- A preliminary evaluation with within-session, cross-session, and cross-user conditions.
- Ablation analyses on sensing channels, temporal windows, and model architecture.

2 Related Work

2.1 Vision-Based and Wrist-Worn Imaging

Vision-based hand pose estimation uses RGB, depth, stereo, infrared, or thermal imagery to infer hand skeletons, meshes, or landmarks. Real-time systems such as MediaPipe Hands have made camera-based hand tracking practical on commodity devices [20]. XR systems similarly use headset-mounted cameras to support pinch, touch, grab, and bare-hand interaction. However, visual tracking can fail under self-occlusion, unusual viewpoints, object occlusion, and poor lighting. Outward-facing cameras also raise bystander and environmental privacy concerns [2, 14].

Wrist-worn imaging systems reduce reliance on external cameras. FingerTrak uses miniature thermal cameras on a wristband for continuous 3D hand pose tracking [6]. Back-Hand-Pose estimates 3D hand pose from dorsal hand deformation observed by a wrist-worn camera [18]. DiscoBand integrates multiple depth sensors into a smartwatch strap for hand, body, and environment tracking [3]. These systems demonstrate the wrist as a useful sensing location, but they still depend on optical observation.

2.2 EMG and Physiological Sensing

Surface electromyography (sEMG) captures muscle activation associated with intended and executed hand motion. NeuroPose and WR-Hand show that EMG-based wearables can estimate continuous 3D hand or finger pose [12, 13]. Large-scale benchmarks such as emg2pose further demonstrate the potential of wrist-sEMG for hand pose estimation across diverse users and movements [15]. However, sEMG can be sensitive to electrode placement, skin contact, perspiration, fatigue, and user anatomy. TacPose instead measures mechanical deformation, avoiding electrodes and electrical skin-contact requirements.

2.3 Force Myography and Deformation Sensing

Force myography (FMG) infers hand activity from pressure changes caused by muscle, tendon, and tissue deformation. FMG has been used for gesture classification, prosthetic control, orthosis control, and finger angle estimation [5, 7, 8, 16, 19]. Compared with EMG, FMG can be mechanically simpler and does not require conductive electrodes. However, it is sensitive to strap tightness, sensor placement, limb posture, and user-specific anatomy.

TacPose can be interpreted as a dense tactile-array version of wrist FMG. Instead of sparse scalar pressure sensors, it uses a 4×6 triaxial tactile array to capture spatial deformation patterns and compare normal-only versus full triaxial sensing.

2.4 Acoustic, RF, and Smartwatch-Based Sensing

Other wrist-worn systems use RF, acoustic, impedance, or smartwatch hardware. EtherPose estimates continuous hand pose from wrist-worn antenna impedance sensing [9]. EchoWrist uses active acoustic sensing for hand pose tracking and hand-object interaction recognition [11]. WatchHand uses off-the-shelf smartwatch speakers and microphones for continuous 3D hand pose tracking [10]. These systems highlight the broader shift toward always-available hand sensing from compact wrist-worn form factors. TacPose contributes a tactile alternative that directly measures local mechanical deformation.

3 System Design

3.1 Hardware Prototype

TacPose uses a XELA Robotics uSkin 4×6 triaxial tactile sensor array mounted on the inner wrist using a commercial wristband and 3D-printed connector. Each of the 24 taxels reports three force-related components: two tangential axes and one normal axis. The sensor is connected to a controller board, powered by a 5 V supply, and streamed to a PC for data collection and inference.

Table 1: Key specifications of the tactile sensor used in TacPose.

Specification	Value
Sensing principle	Hall-effect-based tactile sensing
Taxel layout	4×6 array
Number of taxels	24
Force axes	3-axis force measurement per taxel
Sampling frequency	Up to 83 Hz
Microcontroller voltage	5 V
Dimensions without cable	$30.60 \times 50.60 \times 4.87$ mm
Taxel pitch	7.25 mm

The wrist location is attractive because it is socially accepted, mechanically informative, and does not cover the palm or fingers. Finger movement creates tendon glide and tissue deformation around the wrist, which can be captured as distributed tactile patterns. Unlike gloves or fingertip sensors, the wristband preserves natural touch and object manipulation.

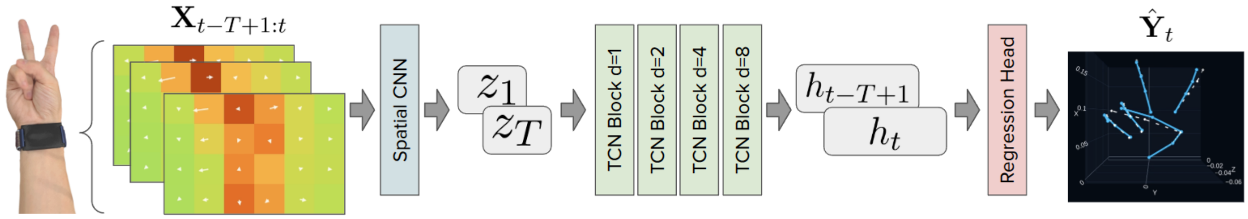

Figure 2: TacPose model architecture. Each tactile frame is encoded spatially, short-term temporal dynamics are modeled with a residual dilated TCN, and the final representation is regressed to wrist-relative 3D hand landmarks.

3.2 Sensing Principle

TacPose is based on the hypothesis that hand and finger motion produces repeatable deformation patterns at the inner wrist. When the user changes hand pose, local pressure and shear patterns change as tendons, skin, and soft tissue move under the sensor. The model learns a mapping from tactile history to hand landmarks:

$$\hat{\mathrm{Y}}_t = f(\mathrm{X}_{t-T+1:t}), \tag{1}$$

where $\mathrm{X}_{t-T+1:t} \in \mathbb{R}^{T\times 4\times 6\times C}$ is a temporal window of tactile measurements and $\hat{\mathrm{Y}}_t \in \mathbb{R}^{J\times 3}$ is the predicted 3D hand pose.

3.3 Target Representation

TacPose predicts wrist-relative 3D landmarks rather than absolute world coordinates:

$$\mathrm{Y}^{rel}_{t,j} = \mathrm{Y}_{t,j} - \mathrm{Y}_{t,wrist}. \tag{2}$$

This removes global hand translation and focuses the learning problem on hand configuration, which is more appropriate for a wrist-mounted tactile sensor.

4 Model

4.1 Spatial-Temporal Tactile Encoder

TacPose formulates wrist-tactile hand pose estimation as a sequence-to-pose regression problem. Given a temporal window of tactile measurements, the model predicts the current wrist-relative 3D hand pose. Each tactile frame is represented as $\mathrm{X}_t \in \mathbb{R}^{4\times 6\times C}$, where $C = 3$ for full triaxial input and $C = 1$ for normal-force-only input.

The model combines a spatial tactile encoder with a temporal convolutional network (TCN). Each frame is reshaped into a $C\times 4\times 6$ tensor and processed by a lightweight 2D convolutional encoder with 32, 64, and 64 channels. Global average pooling and a linear projection produce a 64-dimensional tactile embedding for each frame. This spatial encoder preserves local taxel structure, enabling the model to learn deformation patterns across neighboring sensing points.

The sequence of tactile embeddings is passed through four residual dilated TCN blocks with dilation factors 1, 2, 4, and 8. This allows the model to capture short-term tactile dynamics while keeping the architecture compact. The final timestep representation is then mapped by a regression head to wrist-relative 3D hand landmarks. By combining spatial tactile encoding with temporal modeling, TacPose captures both local deformation patterns and their evolution during hand motion.

4.2 Loss Functions

We train the main model using landmark regression loss:

$$\mathcal{L}_{pose} = \frac{1}{J}\sum_{j=1}^{J} \|\hat{\mathbf{p}}_j - \mathbf{p}_j\|_2^2. \tag{3}$$

We evaluate design choices through ablations on input channels, temporal windows, and model architecture.

5 Data Collection and Evaluation

5.1 Participants and Apparatus

We conducted a preliminary data collection study with five participants to evaluate cross-user generalization. Participants wore the TacPose wristband on the inner wrist. Reference hand pose was collected using a Meta Quest headset with Hand Tracking Streamer, an open-source Meta Quest application for streaming hand tracking landmarks [17]. The resulting dataset contains synchronized tactile frames and 3D hand landmark targets.

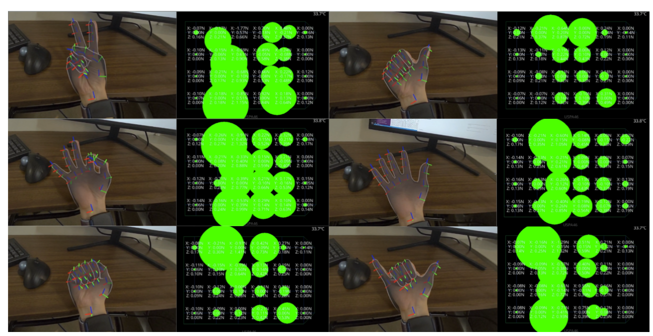

Figure 3: Examples of recorded data pairs. The user's hand and landmarks are visible via pass-through in the Meta Quest 3 headset.

5.2 Procedure

Each participant completed a three-minute random hand movement session. Participants freely moved their fingers and hands through a broad range of configurations, including opening, closing, partial grasps, pinches, individual finger motion, and transitions between poses. This protocol was chosen because TacPose targets continuous pose regression rather than discrete gesture classification.

A single user conducted 10 sessions with varying sensor placements and wristband tension to evaluate cross-session generalization. A single user conducted a 15-minute session for within-session evaluation.

5.3 Dataset Construction

Tactile and hand-pose streams were timestamped and aligned by nearest-neighbor matching within a maximum allowable temporal gap. Tactile data were resampled to a fixed sampling rate and converted into temporal windows. For each window, the target was the wrist-relative hand pose at the final timestep.

5.4 Metrics and Evaluation Conditions

The primary metric is mean per-joint position error (MPJPE), reported in millimeters:

$$\mathrm{MPJPE} = \frac{1}{J} \sum_{j=1}^{J} \|\hat{\mathbf{p}}_j - \mathbf{p}_j\|_2. \tag{4}$$

We report mean MPJPE, MPJPE standard deviation, and median MPJPE. Three evaluation conditions are used:

- **Within-session**: train and test data are from the same recording session.
- **Cross-session**: train and test data are from different sessions.
- **Cross-user**: train and test data are separated by participant.

6 Results

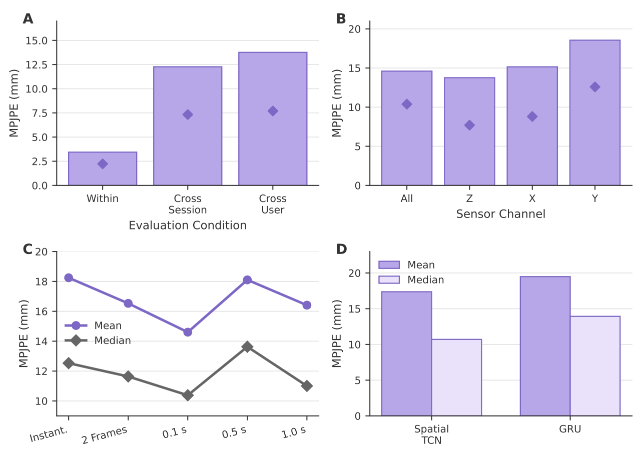

Figure 4: Summary of TacPose evaluation results. (A) Generalization performance across within-session, cross-session, and cross-user settings. (B) Sensor channel ablation under cross-user evaluation. (C) Temporal-window ablation showing the benefit of short tactile history. (D) Architecture ablation comparing the Spatial-TCN model with a GRU baseline.

6.1 Generalization Performance

Within-session evaluation achieved the lowest error, with mean MPJPE of 3.44 mm. Cross-session error increased to 12.27 mm, and cross-user error reached 13.76 mm. This indicates that TacPose captures strong pose-related tactile information under consistent wearing conditions, but generalization is affected by sensor placement, strap tightness, and pressure baseline.

Table 2: Generalization performance across evaluation conditions. Values are reported in millimeters.

Condition	Mean MPJPE	Std.	Median MPJPE
Within-session	3.44	4.67	2.23
Cross-session	12.27	14.56	7.330
Cross-user	13.76	17.15	7.698

6.2 Sensor Channel Ablation

In cross-user evaluation, normal-force-only input outperformed full triaxial input, especially in median MPJPE. This suggests that normal pressure distribution provides a more robust cross-user cue, while tangential channels may be more sensitive to user-specific anatomy, skin friction, and sensor orientation.

Table 3: Sensor channel ablation under cross-user evaluation. Values are reported in millimeters.

Sensor channel	Mean MPJPE	Std.	Median MPJPE
All axes	14.60	15.35	10.38
Z-only	13.76	17.15	7.698
X-only	15.15	19.37	8.804
Y-only	18.56	20.65	12.59

Across conditions, all-axis input slightly improved mean MPJPE in within-session and cross-session evaluation, but reduced cross-user robustness. This indicates that shear channels can contain useful pose information when the sensor-user relationship is stable, but may introduce variability across users.

Table 4: Normal-force-only and full triaxial input across evaluation conditions. Values are reported in millimeters.

Condition / Channel	Mean MPJPE	Std.	Median MPJPE
Within-session / Z	3.44	4.67	2.23
Within-session / All	3.28	4.44	2.11
Cross-session / Z	12.27	14.56	7.330
Cross-session / All	11.75	12.74	7.740
Cross-user / Z	13.76	17.15	7.698
Cross-user / All	14.60	15.35	10.38

6.3 Architecture and Temporal Window Ablations

The Spatial-TCN model achieved lower mean and median MPJPE than the GRU baseline. This supports preserving tactile grid structure before temporal modeling. Temporal-window ablation showed that instantaneous input underperformed short temporal windows. The best mean MPJPE was obtained with a 0.1 s window, while longer windows did not monotonically improve performance.

Table 5: Architecture ablation. Values are reported in millimeters.

Model	Mean MPJPE	Std.	Median MPJPE
Spatial-TCN	17.37	19.79	10.70
GRU	19.49	20.31	13.93

Table 6: Temporal window ablation. Values are reported in millimeters.

Window	Mean MPJPE	Std.	Median MPJPE
Instantaneous	18.25	19.90	12.53
2 frames	16.53	17.45	11.64
0.1 s	14.60	15.35	10.38
0.5 s	18.10	18.38	13.63
1.0 s	16.41	18.35	11.01

7 Discussion

7.1 Feasibility and Generalization

The results support the feasibility of wrist-tactile hand pose estimation. A compact tactile sensor worn at the inner wrist can encode meaningful information about continuous hand pose without imaging the hand, covering the fingers, or using electrodes. The strong within-session result suggests that the tactile signal is informative when wearing conditions are stable.

The larger cross-session and cross-user errors show that practical deployment depends on placement robustness. Sensor rotation, strap tension, local pressure baseline, and skin-device coupling can change the tactile-to-pose mapping. Future systems should address these factors using mechanical alignment, calibration poses, baseline normalization, sensor-shift augmentation, and personalized adaptation.

7.2 Normal and Shear Sensing Tradeoffs

The channel ablations show that normal and tangential sensing provide different tradeoffs. Full triaxial input is helpful when the sensor-user relationship is stable, but normal-force-only sensing generalizes better across users. This suggests that shear signals may encode useful directional deformation, but are more sensitive to placement, skin friction, and anatomy. Larger datasets and channel-normalization strategies are needed to determine whether shear information can be exploited robustly.

7.3 Implications for HCI and Robotics

For HCI, TacPose points toward continuous hand input that is low-obtrusion and relatively privacy-preserving. Because the sensor does not capture images of the user or environment, it can complement vision-based hand tracking in contexts where cameras are unreliable or socially undesirable. Because the palm and fingers remain uncovered, TacPose is compatible with object manipulation and everyday touch.

For robotics, TacPose may support teleoperation, learning from demonstration, and contact-rich manipulation interfaces. A wrist-worn tactile estimator could complement global hand or wrist tracking from VR systems by providing local finger configuration during occlusion. Future extensions may estimate not only pose, but also grasp effort, contact state, or intended stiffness, which are critical for dexterous robot control.

8 Limitations

This study is preliminary. The dataset contains five participants and short controlled recording sessions, so it cannot support strong population-level claims. The protocol uses free hand motion rather than long-term daily use or natural object manipulation. The current system is evaluated primarily using MPJPE, which does not fully capture interaction quality, latency, task success, grasp accuracy, or user experience. Finally, supervised training requires paired hand-pose labels from an external tracking system.

9 Future Work

Future work should collect larger datasets with more participants, repeated donning sessions, diverse wrist postures, arm motion, and object interaction. Modeling work should address placement variability through augmentation, calibration, domain adaptation, and spatial alignment mechanisms. Multimodal fusion with VR hand tracking, IMUs, EMG, acoustic sensing, or smartwatch signals could improve robustness. Task-level studies in XR interaction, accessibility, rehabilitation, robot gripper control, and dexterous teleoperation are needed to evaluate whether offline pose accuracy translates into practical benefit.

10 Conclusion

TacPose demonstrates that a compact wrist-worn triaxial tactile sensor can estimate continuous hand pose from local deformation at the inner wrist. In a preliminary five-participant study, TacPose achieved 3.44 mm within-session mean MPJPE, with higher cross-session and cross-user errors that reveal placement and pressure variability as central challenges. Ablations show that normal pressure is the most robust cross-user cue, while spatial-temporal modeling and short tactile history improve performance. Wrist-tactile sensing offers a complementary path toward continuous, wearable, and privacy-preserving hand sensing for HCI, XR, rehabilitation, robot teleoperation, and learning from human demonstration.

Acknowledgments

This work was conducted as an undergraduate thesis project at Waseda University. Special thanks to Prof. Shigeki Sugano, Yushi Wang, Tamon Miyake, and members of Sugano Lab for their guidance, feedback, and support.

References

[1] Weiya Chen, Chenchen Yu, Chen Tu, Zhiyuan Lyu, Jie Tang, Shiqi Ou, Yan Fu, and Zhidong Xue. 2020. A Survey on Hand Pose Estimation with Wearable Sensors and Computer-Vision-Based Methods. *Sensors* 20, 4 (2020), 1074. doi:10.3390/s20041074

[2] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. 2014. In Situ with Bystanders of Augmented Reality Glasses: Perspectives on Recording and Privacy-Mediating Technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. Association for Computing Machinery, 2377–2386. doi:10.1145/2556288.2557352

[3] Nathan Devrio and Chris Harrison. 2022. DiscoBand: Multiview Depth-Sensing Smartwatch Strap for Hand, Body and Environment Tracking. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, Article 56, 13 pages. doi:10.1145/3526113.3545634

[4] Bardia Doosti. 2019. Hand Pose Estimation: A Survey. *arXiv preprint arXiv:1903.01013* (2019). doi:10.48550/arXiv.1903.01013

[5] Jeremias Gantenbein, Jan Dittli, Jan T. Meyer, Roger Gassert, and Olivier Lambercy. 2023. Feasibility of Force Myography for the Direct Control of an Assistive Robotic Hand Orthosis in Non-Impaired Individuals. *Journal of NeuroEngineering and Rehabilitation* 20, Article 101 (2023). doi:10.1186/s12984-023-01222-8

[6] Fang Hu, Peng He, Songlin Xu, Yin Li, and Cheng Zhang. 2020. FingerTrak: Continuous 3D Hand Pose Tracking by Deep Learning Hand Silhouettes Captured by Miniature Thermal Cameras on Wrist. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2, Article 71 (2020), 24 pages. doi:10.1145/3397306

[7] Tai Inui, Tamon Miyake, Yushi Wang, and Shigeki Sugano. 2025. TriForce Band: Leveraging Triaxial Tactile Sensing for Wrist Force-Myography Gesture Recognition. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '25)*. Association for Computing Machinery, New York, NY, USA. doi:10.1145/3746058.3758418

[8] Xianta Jiang, Lukas-Karim Merhi, Zhen Gang Xiao, and Carlo Menon. 2017. Exploration of Force Myography and Surface Electromyography in Hand Gesture Classification. *Medical Engineering & Physics* 41 (2017), 63–73. doi:10.1016/j.medengphy.2017.01.015

[9] Daehwa Kim and Chris Harrison. 2022. EtherPose: Continuous Hand Pose Tracking with Wrist-Worn Antenna Impedance Characteristic Sensing. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, Article 37, 12 pages. doi:10.1145/3526113.3545665

[10] Jiwan Kim, Chi-Jung Lee, Hohurn Jung, Tianhong Catherine Yu, Ruidong Zhang, Ian Oakley, and Cheng Zhang. 2026. WatchHand: Enabling Continuous Hand Pose Tracking on Off-the-Shelf Smartwatches. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery. doi:10.1145/3772318.3790932

[11] 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 (CHI '24)*. Association for Computing Machinery, Article 352, 19 pages. doi:10.1145/3613904.3642910

[12] Yang Liu, Chengdong Lin, and Zhenjiang Li. 2021. WR-Hand: Wearable Armband Can Track User's Hand. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3, Article 118 (2021), 27 pages. doi:10.1145/3478112

[13] Yilin Liu, Shijia Zhang, and Mahanth Gowda. 2021. NeuroPose: 3D Hand Pose Tracking Using EMG Wearables. In *Proceedings of the Web Conference 2021 (WWW '21)*. Association for Computing Machinery, 1471–1482. doi:10.1145/3442381.3449890

[14] Joseph O'Hagan, Julie R. Williamson, and Mohamed Khamis. 2023. Privacy-Enhancing Technology and Everyday Augmented Reality: Understanding Bystanders' Varying Needs for Awareness and Consent. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4, Article 177 (2023), 35 pages. doi:10.1145/3569501

[15] Sasha Salter, Richard Warren, Collin Schlager, Adrian Spurr, Shangchen Han, Rohin Bhasin, Yujun Cai, Peter Walkington, Anuoluwapo Bolarinwa, Robert Wang, Nathan Danielson, Josh Merel, Eftychios Pnevmatikakis, and Jesse Marshall. 2024. emg2pose: A Large and Diverse Benchmark for Surface Electromyographic Hand Pose Estimation. In *Advances in Neural Information Processing Systems (NeurIPS Datasets and Benchmarks)*. https://arxiv.org/abs/2412.02725

[16] Peter B. Shull, Shuo Jiang, Yifan Zhu, and Xiaoyan Zhu. 2019. Hand Gesture Recognition and Finger Angle Estimation via Wrist-Worn Modified Barometric Pressure Sensing. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 27, 4 (2019), 724–732. doi:10.1109/TNSRE.2019.2905658

[17] Zhengyang K. Weng. 2026. *Hand Tracking Streamer: Meta Quest VR App for Motion Capture and Teleoperation.* https://github.com/wengmister/hand-tracking-streamer

[18] Erwin Wu, Ye Yuan, Hui-Shyong Yeo, Aaron Quigley, Hideki Koike, and Kris M. Kitani. 2020. Back-Hand-Pose: 3D Hand Pose Estimation for a Wrist-Worn Camera via Dorsum Deformation Network. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20)*. Association for Computing Machinery, 1147–1160. doi:10.1145/3379337.3415897

[19] Zhen Gang Xiao and Carlo Menon. 2019. A Review of Force Myography Research and Development. *Sensors* 19, 20 (2019), 4557. doi:10.3390/s19204557

[20] Fan Zhang, Valentin Bazarevsky, Andrey Vakunov, Andrei Tkachenka, George Sung, Chuo-Ling Chang, and Matthias Grundmann. 2020. MediaPipe Hands: On-Device Real-Time Hand Tracking. *arXiv preprint arXiv:2006.10214* (2020). doi:10.48550/arXiv.2006.10214