

# UltraGlove: Hand Pose Estimation with MEMS-Ultrasonic Sensors

Qiang Zhang  
Princeton University, USA  
qz9238@princeton.com

Yuanqiao Lin  
Princeton University, USA  
yuanqiao@princeton.edu

Yubin Lin  
Princeton University, USA  
yubinlin@princeton.edu

Szymon Rusinkiewicz  
Princeton University, USA  
smr@princeton.edu

## ABSTRACT

Hand tracking is an important aspect of human-computer interaction and has a wide range of applications in extended reality devices. However, current hand motion capture methods suffer from various limitations. For instance, visual hand pose estimation is susceptible to self-occlusion and changes in lighting conditions, while IMU-based tracking gloves experience significant drift and are not resistant to external magnetic field interference. To address these issues, we propose a novel and low-cost hand-tracking glove that utilizes several MEMS-ultrasonic sensors attached to the fingers, to measure the distance matrix among the sensors. Our lightweight deep network then reconstructs the hand pose from the distance matrix. Our experimental results demonstrate that this approach is both accurate, size-agnostic, and robust to external interference. We also show the design logic for the sensor selection, sensor configurations, circuit diagram, as well as model architecture.

## CCS CONCEPTS

• **Human-centered computing** → **Interaction devices**; • **Computing methodologies** → **Motion capture**; *Machine learning*.

## KEYWORDS

Hand Tracking, Data Glove

### ACM Reference Format:

Qiang Zhang, Yuanqiao Lin, Yubin Lin, and Szymon Rusinkiewicz. 2023. UltraGlove: Hand Pose Estimation with MEMS-Ultrasonic Sensors. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)*, December 12–15, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3610548.3618202>

## 1 INTRODUCTION

Hand tracking is an essential human-computer interaction (HCI) technology for a variety of applications, such as virtual training systems, virtual/augmented reality (VR/AR) systems, and robotic dexterous manipulation. For example, hand tracking is widely used in the film-making industry, as hand poses captured from actors are reprojected to animated characters, known as avatars, for better realism. Similarly, hand tracking benefits athletes, since they are able to record their motions and later examine them to refine their technique. Hand tracking also has important implications for the field of human-robot interaction. Humanoid robots with dexterous

hands can adapt to complicated and dangerous scenarios and replace human labor. If captured hand poses and motions are made available, they can either be used in the teleoperation of robots or serve as training data for learned controllers.

Existing hand-tracking systems can be categorized by their sensing mechanisms: vision-based, IMU-based, and stretch-based. Visual tracking directly predicts hand motions from one or more RGB or RGB-D cameras, but such systems are sensitive to self-occlusion and limited to the field of view: they are prone to failure when hands are obscured from the cameras during manipulation. Additionally, background variations, such as insufficient lighting or excessive movements, also interfere with the extraction of hand poses, leading to low accuracy and inconsistent results. Both inertial- and stretch-based hand tracking can be used without the constraint that the hand must remain inside a camera's field of view, but they still have their drawbacks. Inertial-based measurements are taken from a grid of inertial measurement units (IMUs) attached to a glove, but those units are sensitive to external magnetic interference and lack long-term stability due to sensor drift. Additionally, they cannot distinguish between different finger poses with similar orientations, which limits their applicability in certain scenarios. Stretch-based methods rely on stretch sensors attached to the fingers to measure the degree of bending, but they cannot be easily adapted to people with different hand sizes and cannot easily differentiate between open- and closed-finger poses (i.e., distinguishing whether fingers are separated laterally).

Overall, the potential applications of hand-tracking technology are diverse and have far-reaching impacts from entertainment to manufacture. However, existing hand-tracking methods have certain limitations, such as low accuracy, low robustness to external interference, and lack of adaptability to different hand sizes. These drawbacks restrict their applications and hinder the deployment of hand tracking in high-compliance scenarios, such as robot control or remote surgery. Therefore, controllers are still widely used for hand pose commands, which require users to map an intuitive action (i.e. grabbing an object) to an abstract interaction (e.g., pressing a button on the controller). This reduces the naturalness and available degrees of freedom of the interaction, often leading to unintuitive cognitive mapping, time-consuming training, and reduced precision. By developing a low-cost and accurate hand-tracking solution, this paper aims to make this technology more accessible and widely applicable, paving the way for further innovations in the field.

In this paper, we propose a novel and low-cost solution for hand tracking using a glove with multiple micro-electromechanical system (MEMS) ultrasonic sensors positioned throughout the hand. At run time, we measure the pairwise distances among the sensors and reconstruct the hand pose using a lightweight deep neural network. We have conducted extensive experiments to evaluate our method's performance in both mechanical hands with quantitative metrics and in human hands with qualitative metrics. The results

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia*

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0315-7/23/12.

<https://doi.org/10.1145/3610548.3618202>

demonstrate that our approach achieves high accuracy and is robust to interference under challenging scenarios that existing methods cannot handle, thus making it suitable for various HCI devices, virtual training systems, and robotic dexterous manipulation.

Our main contributions are as follows:

- We design and build a glove that integrates multiple MEMS-ultrasonic sensors, obtaining a pairwise distance matrix for each hand pose. We also develop a circuit and implement the corresponding embedded system to read the raw sensor data and measure the distance matrix in real time.
- We propose a lightweight deep neural network model for accurate and real-time 3D hand pose estimation based on the raw sensor data returned from the sensor. We collect a training dataset and evaluate the performance of the proposed model, demonstrating its effectiveness through sim-to-real transfer learning.
- We provide an in-depth analysis of the design philosophy for the raw sensor selection, sensor configurations, microcontroller unit (MCU), and circuit architecture. We also conduct an ablation study on the proposed model to evaluate the contribution of each component.

## 2 RELATED WORK

### 2.1 Visual Hand Pose Estimation

There has been significant progress in hand pose estimation using RGB or RGB-D cameras. For example, marker-based data gloves have been proposed [Han et al. 2018; Wang and Popović 2009], which require colored or optical markers to be attached to the glove and rely on external cameras to estimate pose.

In the absence of markers, some systems have proposed identifying 2D keypoints in images of hands, most recently based on convolutional neural networks that produce keypoint heatmaps [Cai et al. 2020; Iqbal et al. 2018]. Some methods [Cai et al. 2018; Mueller et al. 2018; Spurr et al. 2020; Wang et al. 2020; Zimmermann and Brox 2017] directly predict the 3D skeleton from a single image. For example, [Cai et al. 2018] proposes a weakly-supervised 3D hand pose estimation algorithm from monocular RGB images. With the proven success of transformer network architectures, some papers [Li et al. 2022; Lin et al. 2021a,b] have proposed transformer-based or attention-based networks for hand pose estimation. For example, [Li et al. 2022] uses an attention mechanism to model both pose and shape, exploiting the MANO [Romero et al. 2022] prior.

To provide robustness to occlusion, some works focus on multi-view fusion via triangulation [Simon et al. 2017], post-inference optimization [Han et al. 2020], or latent-feature fusion [He et al. 2020; Iskakov et al. 2019; Remelli et al. 2020]. For example, [Han et al. 2022] proposes a differentiable end-to-end architecture for multi-view camera fusion and temporal fusion to improve performance and robustness. Two-hand reconstruction exacerbates the occlusion problem, and many works have addressed occlusion- and collision-aware two-hand joint pose estimation, such as [Fan et al. 2021; Kim et al. 2021; Moon et al. 2020; Rong et al. 2021; Zhang et al. 2021].

RGB-D cameras provide extra information for hand tracking, relative to simple RGB images. Many papers have proposed deep-learning-based algorithms for single-hand tracking [Moon et al. 2018; Mueller et al. 2019; Oikonomidis et al. 2011; Tang et al. 2014,

2015; Xiong et al. 2019] or two-hand tracking [Kyriazis and Argyros 2014; Mueller et al. 2019; Oikonomidis et al. 2012; Tzionas et al. 2016] from a single depth image. For example, [Tang et al. 2015] shows a new hierarchical sampling optimization method to regress the full pose from a depth image via surrogate energy selection.

### 2.2 IMU-based Data Gloves

Many data gloves use Inertial Measurement Units (IMUs) for hand tracking. Researchers have explored systems based on different numbers of IMUs, ranging from 12 [Hu et al. 2020] to 15 [Fang et al. 2017], 16 [Chang and Chang 2019; Connolly et al. 2017], or 18 [Lin et al. 2018]. There are also many works focusing on full body pose reconstruction via sparse (only 6) IMUs, [Huang et al. 2018; Jiang et al. 2022; Von Marcard et al. 2017]. The major drawback of this IMU-based solution is that the raw sensor is sensitive to external magnetic fields, which can cause measurement drift and requires calibration from time to time. Another disadvantage is that the absolute tracking accuracy is limited due to the need to integrate the output of gyroscopes and accelerometers, which is sensitive to noise and miscalibration.

### 2.3 Stretch-sensor-based Data Gloves

Gloves based on stretch sensors have been explored in the context of gesture recognition, such as [Hammond et al. 2014; Lorussi et al. 2005; O'Connor et al. 2017; Ryu et al. 2018]. However, their classification methods can not decode full continuous hand pose. In contrast, [Chossat et al. 2015; Park et al. 2017] propose using stretch sensors for continuous pose estimation, but there is no qualitative regression accuracy reported. [Glauser et al. 2019] presents a promising approach for obtaining continuous pose using a stretch-sensor-based glove. However, the system can not distinguish the opening and closing state of the palm according to their website video [Glauser 2019], and the fabrication of such a glove is complicated.

### 2.4 Other Sensor Data Gloves

There are other hand-tracking solutions with different sensors. For bend-sensor data gloves ([Ciotti et al. 2016; Shen et al. 2016; Zheng et al. 2016]), the number of degrees of freedom is much less than that of a human hand, and increasing the number of bend sensors leads to a high complexity of glove design and may hinder dexterous movement. Trackers based on electromyography (EMG), such as [Liu et al. 2021], require initialization and calibration for every new user and under different sensor locations. Also, the same hand poses with different forces may lead to completely different EMG signals, severely compromising the algorithm's accuracy. Electronic skin sensor solutions such as [Kim et al. 2022] cannot decode the full hand pose and can only be used for some specific applications. *Digits* [Kim et al. 2012] proposes to estimate 3D hand pose through a wrist-worn IR camera. However, it is sensitive to occlusion and requires the hand to be always inside the field of view, reducing the allowable space of wrist angles by half.

Different from all the hand pose estimation methods mentioned above, we propose a novel data glove via ultrasonic sensors. Some previous works also use ultrasound sensing for motion capture, such as body capture: [Laurijssen et al. 2015; Qi et al. 2014; Sato et al.

2011; Vlasic et al. 2007]. However, their capture quality is relatively low and cannot be replicated to the hand pose scale. There are also some ultrasonic-based hand gesture recognition works, such as [Yang et al. 2018, 2020]. However, their methods can only solve the hand gesture classification task with limited categories and lack continuous motion decoding. However, our data glove can predict the full hand pose in a continuous way: on the low-level sensor side, these ultrasonic sensors measure their absolute distances to other sensors and return the mutual pairwise distance matrix in real time with high refresh rate and high spatial resolution. On the high-level algorithm side, our deep network takes this matrix as the input and predicts the hand pose.

In summary, our ultrasonic-based data glove overcomes the self-occlusion challenges in visual tracking systems, overcomes the low robustness to external magnetic fields of IMU-based systems, and overcomes the opening-closing ambiguity of stretch-based systems, while providing high accuracy.

### 3 GLOVE SYSTEM DESIGN

#### 3.1 MEMS-ultrasonic Sensor Introduction

Traditional ultrasonic distance sensors use piezoelectric crystals to generate and receive high-frequency sound waves. Transducers convert electrical energy into mechanical vibrations, creating sound waves that are detected and converted back to electrical signals. The time delay between transmission and reception is used to calculate the distance between the transmitter and receiver. Traditional sensors are typically large compared with finger size. Moreover, their distance measurement accuracy level is around 10 – 20 mm and they have relatively low beam widths, both of which do not satisfy our system requirements. Please refer to Section 5.2 for more details.


In contrast, ultrasonic sensors based on micro-electromechanical systems (MEMS) are built with micromachining technology and thus are small and highly sensitive. MEMS technology allows for the creation of miniaturized, integrated sensors that can be mass-produced at low cost. Compared with piezoelectric-based sensors, they have a smaller size, lower power consumption, and most importantly, their accuracy level is much higher: we will illustrate how much accuracy they can achieve in the experiment section.

Here we choose the CH101 ultrasonic sensor from TDK, with a 4x4x2mm size. In our application scenario, we need the beam angle of the ultrasonic sensor to be as wide as possible, and the CH101 satisfies this need – in our experiments, its horizontal and vertical beam angles can be as wide as 150 degrees. This is necessary to avoid missing measurements when we attach these sensors to the fingers and measure their pairwise distance matrix.


In our prototype, we attach 7 sensors to the hand as shown in Fig. 1. Subfigure (a) shows a single sensor, (b) demonstrates how they are attached, and (c) shows how the circuit and the sensors are connected.

#### 3.2 Sensor Data Acquisition System

We choose 7 CH-101 sensors that communicate with the SmartSonic development board from TDK using the I2C protocol. During measurement, we cycle through sensors 1 to 7 to select one sensor at a time as the transmitter, while the remaining sensors serve as



**Figure 1: Visualization of sensors and how they are attached to the human hand in our system. From left to right: (a) Single CH101 sensor. (b) The back side of the hand with the sensors attached. (c) The front side of the hand and the embedded system circuit.**



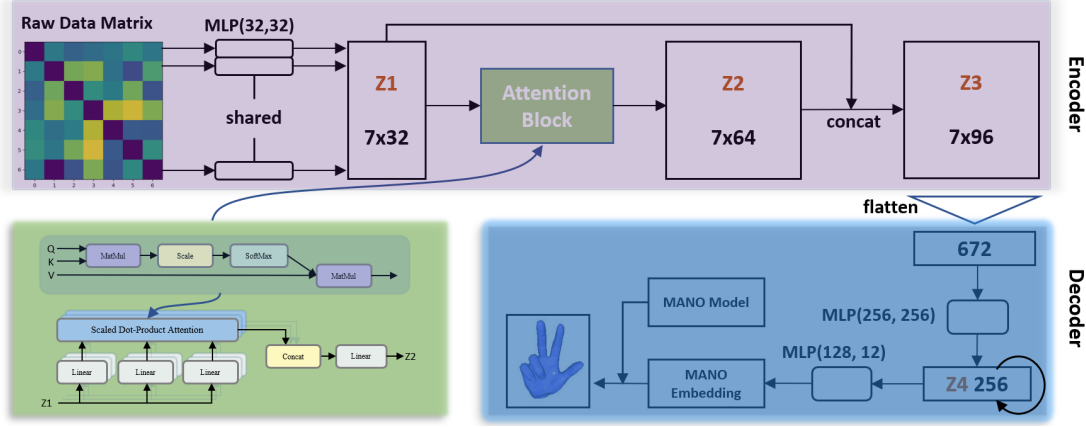
**Figure 2: System-level diagram visualization for data acquisition with the SmartSonic development kit. The development kit provides an I2C bus to all 7 sensors at 1.8V. It also provides enough I/O pins that are internally level-shifted for the 4 ultrasound sensors. However, external level-shifters are used for the I/O pins of the remaining 3 sensors. The board uses a serial link to send sensor readings to the computer.**

receivers. This approach enables us to obtain six distance values simultaneously and create a complete distance matrix within a single cycle. Then the development board relays the raw sensor matrix to the laptop using the serial protocol.

As shown in Fig. 2, the development board only provides enough level-shifted I/O ports for 4 sensors. Therefore, additional off-board level-shifting circuits are used to translate between 3.3V and 1.8V logic. For unidirectional buses, resistor dividers provide adequate performance due to the low acquisition rate in this system. A discrete-part translator from SparkFun (BOB-12009) is used to drive bidirectional pins. However, this system setup does not scale well due to development board limitations. We describe a scalable system that can support more CH-101s with a commercially available MCU in Section 5.4.

#### 3.3 Encoder-Decoder Pose Prediction Model

The low-level embedded system collects the raw pairwise distance matrix from the seven MEMS-based ultrasonic sensors, which is represented as a  $7 \times 7$  matrix. This matrix is then fed into our pose prediction model, which predicts the hand pose represented by 23 joint positions. Our model, shown in Fig. 3, consists of an encoder with an attention mechanism and a decoder that incorporates information from previous frames as well as the MANO hand prior [Romero et al. 2022]. The encoder maps the  $7 \times 7$  distance



**Figure 3: Pose Prediction Model Framework.** Our model consists of an encoder and a decoder. The encoder module takes the raw data matrix as input, then feeds it into an MLP, followed by the attention block, whose output is concatenated with the MLP feature. For the decoder, we first flatten the feature and feed it into an LSTM to incorporate information from past frames and the MANO model to incorporate a hand pose prior.

matrix into a  $7 \times 96$  feature space, and the decoder uses this feature to predict the joint positions. Here we describe them in detail.

*Encoder Module:* For each sensor, we pass its vector of distances to other sensors into an MLP ( $7 \rightarrow 32 \rightarrow 32$ ) model, then stack the results to get a  $7 \times 32$  feature embedding ( $Z_1$ ).

We then use a self-attention module to both provide robustness against missing measurements and extract the graph information among these sensors – intuitively, how these sensor distances formulate the hand pose pattern. To be specific, we use the classical multi-head attention model, built upon scaled-dot-product attention blocks:

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{D_k}}\right)V = AV, \quad (1)$$

where  $Q$ ,  $K$  and  $V$  are the query, key, and value, and  $D_k$  is the dimension of the key. In our model, we use different linear transformations of  $Z_1$  for the query, key, and value in each attention head:

$$\text{head}_i = \text{Attn}(Z_1 W_i^Q, Z_1 W_i^K, Z_1 W_i^V), \quad (2)$$

where  $W_i^Q$ ,  $W_i^K$  and  $W_i^V$  are learnable  $32 \times 64$  matrices. Finally, we concatenate these heads and multiply them with a final linear matrix  $W^O$  to get the  $7 \times 64$  feature embedding  $Z_2$ :

$$Z_2 = \text{MH-Attn}(Q, K, V) = \text{Concat}(\text{head}_1, \text{head}_2, \dots) W^O. \quad (3)$$

We concatenate  $Z_1$  and  $Z_2$  with a skip-connection to get the final encoded feature  $Z_3$  with shape  $7 \times 96$ :

$$Z_3 = \text{Concat}(Z_1, Z_2). \quad (4)$$

*Decoder Module:* We first flatten  $Z_3$  into a one-dimensional vector, which is then followed by another MLP ( $672 \rightarrow 256 \rightarrow 256$ ) to convert it into  $Z_4$ , with size 256:

$$Z_4 = \text{MLP}(\text{Flatten}(Z_3)). \quad (5)$$

To aggregate information from previous time steps, we use an LSTM model with the hidden dimension the same as the input dimension 256. The LSTM cell takes as input a sub-sequence of feature

vectors  $Z_4^1, Z_4^2, \dots, Z_4^T$ , where  $T$  is the length of the sub-sequence. For each sub-sequence, the LSTM processes each feature vector  $Z_4^i$  in order and updates its internal state. After the last feature vector in the sub-sequence is processed, the final hidden state of the LSTM is used as a summary and represents the aggregated information from the previous five timesteps:

$$F = \text{LSTM}(Z_4^1, Z_4^2, \dots, Z_4^T), T = 5. \quad (6)$$

The MANO (Model for Articulated Hands) model [Romero et al. 2022] is a parametric 3D hand model that represents the human hand as a set of articulated bones, joints, and skin. It can be used to generate 3D hand poses from a set of input parameters. During the MANO hand training phase, a large amount of hand pose data is collected and subjected to PCA analysis, resulting in a set of principal component vectors. These principal components represent the patterns of variation in hand poses (joint angles). By adjusting the weights assigned to these principal components, different joint angles can be generated.

The weight parameter dimension of this MANO model is a hyper-parameter, and here we set it as 12. We feed the feature information from the temporal LSTM module described above into an MLP ( $256 \rightarrow 128 \rightarrow 12$ ) and pass its output to the MANO hand model:

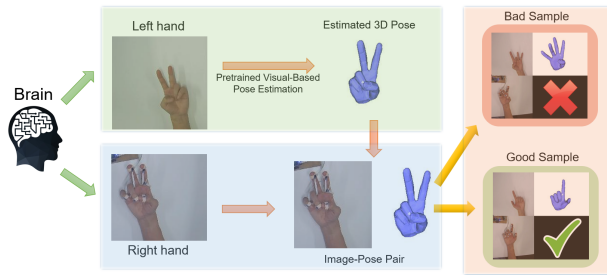
$$J = \{J^1, J^2, \dots, J^n\} = \text{MANO}(\text{MLP}(F)), \quad (7)$$

where the  $J^i$  are the parameters of the joints and  $n$  is the number of degrees of freedom: 23 for the human hand, though Sec. 4.4 presents an experiment with a mechanical hand with  $n = 5$  DOF.

Our entire model – encoder and decoder – is trained end-to-end with L2 loss. Training takes 1-2 hours on a single Nvidia 3080Ti.

### 3.4 Sim-to-real Training Pipeline

To achieve the best possible performance in hand pose estimation, we adopt a sim-to-real transfer training pipeline. This pipeline involves several steps. First, we consider the motion of a simulated hand, and compute sensor positions relative to that hand in the same configuration as we use in our real-world physical setup.



**Figure 4: Dataset collection system visualization.** In this figure, we visualize how to collect the raw image and obtain its (pseudo-)ground truth pose via left-right-hand synchronization, a pretrained hand pose estimation model, and the human filtering process. The visualization of hand poses is derived from [Romero et al. 2022].

This ensures that the simulated data captures the same physical interactions between the hand and the sensors as in the real world.

Secondly, we position the simulated hand using sequences of poses from the InterHand2.6m dataset [Moon et al. 2020], and compute the distances between the simulated sensors. We augment the distance measurements by adding noise and randomly mask measurements to simulate missing data. This allows us to generate a large amount of labeled training data in a controlled and reproducible way that nevertheless resists overfitting and accommodates a wide range of hand sizes and measurement imperfections.

Next, we train a sequential pose prediction model using the simulated dataset. This model takes the sequence of hand poses as input and predicts the next pose in the sequence. By training on the simulated data, the model learns to generalize well to variations in hand shape and movement.

Finally, we adapt the model to the real-world domain by fine-tuning it on a real dataset (described below). This sim-to-real transfer training pipeline has been shown to be effective in improving the performance of hand pose estimation models, especially in scenarios where large amounts of labeled real data are not available.

### 3.5 Dataset and Pseudo-Ground Truth Collection

*System Setup:* We obtain a dataset, used for both fine-tuning our model and qualitative evaluation, consisting of raw sensor distance data synchronized with a pseudo-ground truth obtained using a vision system. To account for processing delays that can occur when collecting data in a single process, we adopt a multi-process collection approach, in which one process handles the raw data and another computes the ground truth from a camera video stream. These two processes each use their own data buffers.

*Pseudo-Ground Truth Extraction:* It is not easy to manually label each hand pose from scratch, nor to perform visual tracking on a hand with our sensors attached to it. Nevertheless, we find that our training procedure benefits from a small amount of real-world fine-tuning, even with imperfect “ground truth.” We therefore choose a procedure that sacrifices accuracy in the ground truth in return for ease of acquisition. Specifically, we mount our sensors to a person’s

right hand and ask them to perform a set of hand poses while mirroring what they do with their left hand, using the latter (with a mirror reflection) for our vision-based system [Li et al. 2022] (see Fig 4). This introduces the possibility of several kinds of error, most notably that the vision system fails or that the left- and right-hand poses are not synchronized. To combat the former, we manually screen the output of the vision system and remove any outputs that do not appear to match the video frames. For the latter, we remove frames in which any finger’s position differs by more than 4mm between the pre-trained estimated pose and the vision pose. In all, we end up filtering out approximately 15% of frames, which still leaves us with a dataset that is useful for fine-tuning.

*Hand Position and Orientation Normalization:* Since our system does not predict global hand position and orientation, we remove this information from the dataset by normalizing each hand pose. We first shift the whole hand such that the wrist point lies at the origin, then rotate the hand such that the root of the middle finger is located on the Z-axis and the root of the index finger is located on the X-Z plane.

## 4 EXPERIMENTAL VERIFICATION AND APPLICATIONS

### 4.1 Dataset Statistics

We use the InterHand2.6m dataset [Moon et al. 2020] for pre-training. This is a large-scale hand pose estimation dataset containing over 2.6 million hand images with corresponding 3D hand joint annotations. Although there are 2.6m images, they represent multiple views for each of around 46k hand poses. We use all of these 46k hand poses as the simulation dataset. Some pose samples are visualized in Fig. 11a.

Our real-world dataset contains around 5000 frames, each containing a raw distance matrix and the hand pose represented as 23 joint positions. For the distance matrix, when one sensor misses the signal sent from another sensor, we mark the value as -1. This missing data accounts for less than 1% of the whole dataset. Fig. 11b shows sample frames from our real-world dataset, while Fig. 11c visualizes the corresponding distance measurements.

### 4.2 Raw Sensor Accuracy Analysis

Here we describe an experiment to evaluate the accuracy of the raw sensor data. As shown in Fig. 6, we attach three sensors (named A, B, and C) at the vertices of an equilateral triangle on a rotating platform. There is also another sensor D attached to the nearby box, and the box is always fixed. We then rotate the platform and collect the sensor distances between D and A, D and B, as well as D and C. We compute the analytical position for sensor D based on its distances to other sensors, in a coordinate system centered at C and with the direction from B to A as the x-axis and the vertical direction as the z-axis. We thus expect the computed coordinates of D to lie in a circle, as the platform is rotated.

The point positions projected to the X-Y plane are shown in Fig. 7a, and we can fit a circle to these points as shown in Fig. 7b. It is qualitatively clear that all the points lie close to the circle. Quantitatively, the average localization error is 0.65mm, and most errors are below 1.2mm, as shown in the histogram in Fig. 7c.

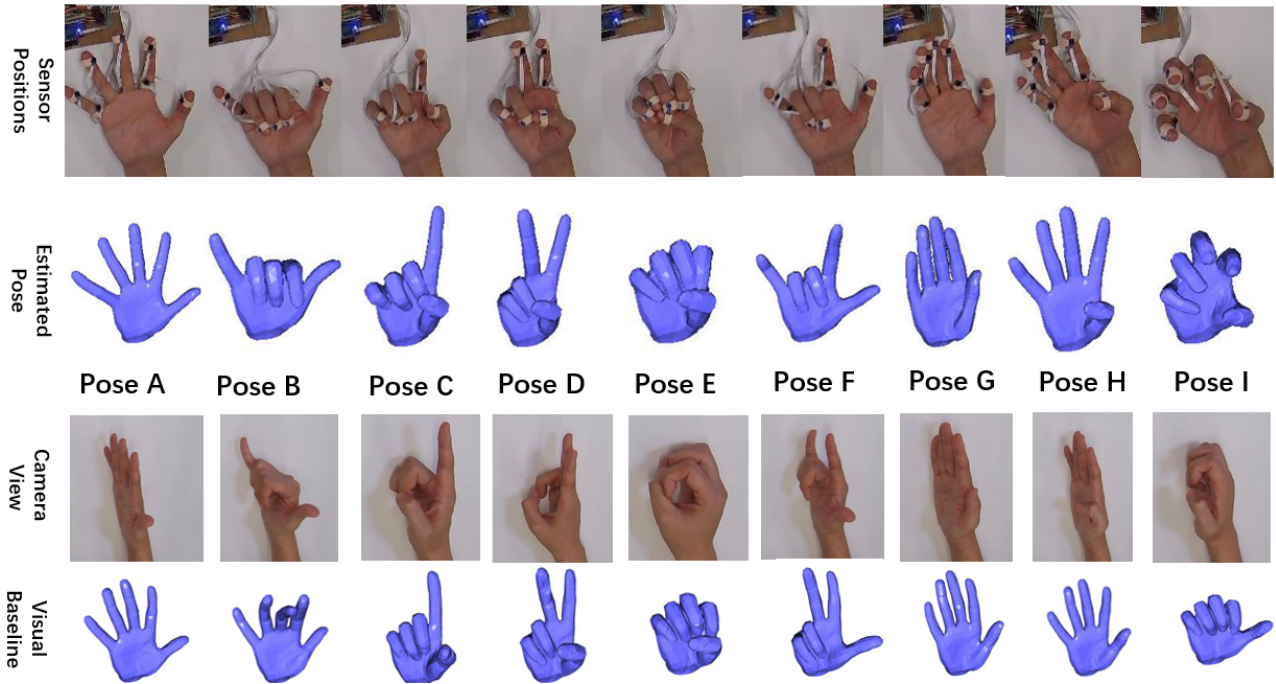


Figure 5: Qualitative Performance Evaluation. From top to bottom, we show a hand with our sensors attached, the estimated hand pose, the mirrored camera view image for the same pose, and the vision baseline results. Our method qualitatively outperforms the vision baseline, especially in the presence of occlusion.




Figure 6: Evaluation of localization accuracy. Sensors A, B, and C are attached to the rotating platform, while sensor D is fixed on the nearby box. (a) and (b) show the setup at two different angles of rotation.





Figure 7: The sensor D is localized relative to a coordinate system defined by A, B, and C. (a) Scatterplot of the coordinates of D as the platform is rotated. (b) Best-fit circle, showing the qualitative accuracy of localization. (c) Histogram of distances between points and the best-fit circle.

### 4.3 Qualitative Evaluation on Human Hand Poses


We train our model on real human hands using the sim-to-real pipeline described above, and use it to visualize the qualitative performance of our system. Fig. 5 shows a series of poses, with an image of the hand wearing the sensors at top and our predicted hand pose (rendered in the Open3D engine, with the MANO hand representation [Romero et al. 2022]) below. We compare our results to a vision baseline [Li et al. 2022], demonstrating the frequent inaccuracy in the latter, especially for poses that incorporate significant occlusion with respect to the camera.

*Size-agnostic Hand Pose Estimation:* Our pose estimation model is designed to be agnostic to hand size, given that it is trained on a synthetic dataset containing hands of multiple shapes. Figure 12 evaluates hand poses estimated for multiple individuals with hands of different sizes and shapes. The top row shows results for the same person on which the model was fine-tuned, while the other rows demonstrate good adaptability to individuals with smaller hands. The hands from top to bottom have lengths of 21.4 cm, 19.2 cm, and 15.5 cm, and widths of 8.1 cm, 7.7 cm, and 7.0 cm. Given the large variation in hand sizes, the system generalizes well though not perfectly.

*Pseudo Ground-truth Quantitative Comparisons:* The average distance between our predicted joint positions and the pseudo ground-truth joint positions, as measured on the test set, is 1.09 cm. Although we cannot determine how much of the error comes from



**Figure 8: (a) Mechanical hand with five degrees of freedom, used for quantitative evaluation. (b) Hand with sensors attached, in the same configuration as that for the human hand.**



**Figure 9: Results of accuracy experiment using mechanical hand. The graphs show results for the thumb, index, middle, ring, and pinky fingers, respectively. The horizontal axis is the ground-truth servo command, while the vertical axis represents the inferred servo command.**

the pseudo ground truth error vs. our proposed system framework/model, it demonstrates that the error level of our system is relatively low.

*Nearest Neighbor Baseline:* We also provide the results for nearest neighbor baseline: during the inference stage, we measure the similarity between the real-time distance matrix from the embedded system and the dataset distance matrices, and select the hand pose with the maximum similarity to represent the model prediction. Here the similarity is defined as the cosine distance between the two 49-d vectors (flattened from  $7 \times 7$  matrix). It turns out that the quantitative result for this baseline is 3.64cm, which is much higher than ours and demonstrates the effectiveness of our model.

#### 4.4 Performance Evaluation Using a Mechanical Hand

To quantitatively evaluate our system, we conducted experiments using a mechanical hand (Fig. 8) with five degrees of freedom. Each finger is controlled by a separate servo motor, and we can therefore define a 5-dimensional space of poses parameterized by the 5 servo commands.

To accommodate the difference between the 23-dimensional human hand model and the 5-DOF mechanical hand, we remove the MANO stage from the original model and replace it with an MLP ( $256 \rightarrow 128 \rightarrow 5$ ) model, which directly regresses the five-finger command signals.

Our dataset consists of 30,000 frames, each including a raw data distance matrix and the corresponding hand servo commands representing the hand pose. Fig. 9 shows results from a testing dataset, comparing actual and predicted poses. In each graph, the horizontal axis denotes the ground truth pose value (normalized to the range  $[-0.5, 0.5]$ ), while the vertical axis represents our prediction result. The mean error was found to be 0.0163, demonstrating that our model can achieve excellent performance.

## 5 DISCUSSION

### 5.1 Ablation Study of Our Model

Here we provide an ablation study for the model we designed for the hand pose estimation to illustrate the effectiveness of each module. As shown in Table 1, these three ablation study experiments represent removing the sequential module, attention module and skip connection. We find that all three of these modules contribute to the final performance of our system.


We also provide an ablation study for the fine-tuning sim2real stage. Recalling the pseudo ground-truth quantitative metric, fine-tuning reduced the error from 1.45 cm to 1.09 cm, thus demonstrating nontrivial improvement. This effectiveness comes from differences between the simulator and the real-world dataset: 1. Slightly different sensor positions between the simulator and the real world. 2. Different coverage of hand poses between simulation and the real world. 3. Lack of physical effects such as ultrasound reflection or transmission through occluders in the simulation. It would be possible but not easy to add these to our simulator, since they would require significant calibration against the real world.

**Table 1: Ablation study of components in our pose prediction model. The full model with sequential processing, attention mechanism, and skip connection provides the highest accuracy.**

	w/o seq.	w/o atten.	w/o skip	full
L2 loss	0.0196	0.0215	0.0207	<b>0.0163</b>

### 5.2 Alternative Ultrasonic Sensor Selection

We also experimented with another, more traditional type of ultrasonic sensor that is based on the piezoelectric effect. As shown in Fig 10a, the upper image is the individual sensor and the lower one is the sensor without its outer casing. We have removed the casing to enable transmission and reception across a wider beam width. However, the improvement in beam angle was not sufficient, and we found that we could not receive signals emitted from beyond approximately 60 degrees. Consequently, we designed a dodecahedron-shaped support frame as shown in Fig 10b and 3D-printed this support frame, which allows the placement of 12 sensors. This dodecahedral sensor array is omnidirectional for both transmitting and receiving ultrasound waves. Since the array can be driven directly from the MCU ports without an intermediate translator or driver, a refresh rate of up to 500Hz was achieved.



**Figure 10: Dodecahedral design for omnidirectional sensing using low-cost, narrow-beam piezoelectric ultrasonic transducers. (a) Individual sensor with and without its casing. (b) 3D CAD design for the support frame. (c) Assembled sensor array.**

However, despite the potential benefits in refresh speed, omnidirectionality, and cost, we ended up not using this sensor configuration based on two drawbacks. First, the assembled array is too large, with a radius above 15mm. While this may not be a problem for other tracking applications, such as body pose capture, the size rendered it unfit for attachment to fingers. Secondly, the distance resolution is relatively low and the noise level is higher than the MEMS-ultrasonic sensors due to low ultrasound frequency.

### 5.3 Different Sensor Configurations Design

We also experimented with several different sensor configurations in our system. The number of sensors can be varied from 5 to 8 or more. Here we analyze the trade-offs between different sensor configurations. The minimum number of sensors is 5, and we attach them at the fingertips. The number of pairwise distances available from these sensors is only 10, which is lower than the number of degrees of freedom of a human hand. Moreover, this minimal number of sensors provides very little robustness to missing measurements due to occlusion. If we add one more sensor, the best place to put it is on the wrist. However, this sensor attachment method is not stable due to the movement of the wrist. For seven sensors, we place the two additional sensors at the root of the index and little finger for optimum performance. The performance gain plateaus with additional sensors beyond seven regardless of the sensor placement.

Quantitatively, we obtain simulated mean errors of 1.24 cm, 1.07 cm, 0.85 cm, and 0.82 cm for 5-8 sensors, respectively. The results match with the real experiment results. Based on these analyses and experiments, we motivate our conclusion that a 7-sensor setup offers the best tradeoff between accuracy and being easy to wear.

### 5.4 Different MCU and Embedded System Design

In this section, we propose a general framework for expanding the number of supported sensors. Though developed mainly for hand motion capture, this system has the potential to scale to a variety of motion capture settings, such as wrist, arm, torso, or whole-body movement. Therefore, sensor scalability is an essential issue for the universality of the algorithm. The development kit used to demonstrate the hand-tracking system in this work only supports up to 8 sensor nodes. To achieve a wider range of motion capture applications, more sensors are needed to maintain the spatial and temporal resolution of the dataset. We present a solution for scalable deployment for higher numbers of sensor nodes in the supplemental material. Dramatically increasing the number of sensors will also require solutions such as frequency multiplexing to avoid linear growth in the time of each measurement cycle.

## 6 CONCLUSION

We propose a novel hand motion capture glove based on MEMS-ultrasonic sensors. Our work represents a non-trivial improvement in the field of hand-tracking, as it addresses the limitations of existing solutions and provides a practical and low-cost alternative for accurate and robust hand pose estimation. The proposed design and methodology can be applied to various applications such as virtual reality, human-computer interaction, and dexterous robot manipulation. The main limitations of our system include our real-world fine-tuning, which relies on mirrored hand motion, as well as the fact that our accuracy degrades when an object is held in the hand, since this blocks the propagation of ultrasound waves. However, this can be solved by attaching more sensors, with some on the front of the hand and others on the back side of the fingers – this is a possible direction for future work. Another avenue of future work includes extending the current framework to full human body pose estimation.



## REFERENCES

- Yujun Cai, Lihao Ge, Jianfei Cai, Nadia Magnenat Thalmann, and Junsong Yuan. 2020. 3D hand pose estimation using synthetic data and weakly labeled RGB images. *IEEE transactions on pattern analysis and machine intelligence* 43, 11 (2020), 3739–3753.
- Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. 2018. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *Proceedings of the European conference on computer vision (ECCV)*. 666–682.
- Hsien-Ting Chang and Jen-Yuan Chang. 2019. Sensor glove based on novel inertial sensor fusion control algorithm for 3-D real-time hand gestures measurements. *IEEE Transactions on Industrial Electronics* 67, 1 (2019), 658–666.
- Jean-Baptiste Chossat, Yiwei Tao, Vincent Duchaine, and Yong-Lae Park. 2015. Wearable soft artificial skin for hand motion detection with embedded microfluidic strain sensing. In *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2568–2573.
- Simone Ciotti, Edoardo Battaglia, Nicola Carbonaro, Antonio Bicchi, Alessandro Tognetti, and Matteo Bianchi. 2016. A synergy-based optimally designed sensing glove for functional grasp recognition. *Sensors* 16, 6 (2016), 811.
- James Connolly, Joan Condell, Brendan O'Flynn, Javier Torres Sanchez, and Philip Gardiner. 2017. IMU sensor-based electronic goniometric glove for clinical finger movement analysis. *IEEE Sensors Journal* 18, 3 (2017), 1273–1281.
- Zicong Fan, Adrian Spurr, Muhammed Kocabas, Siyu Tang, Michael J Black, and Otmar Hilliges. 2021. Learning to disambiguate strongly interacting hands via probabilistic per-pixel part segmentation. In *2021 International Conference on 3D Vision (3DV)*. IEEE, 1–10.
- Bin Fang, Fuchun Sun, Huaping Liu, and Di Guo. 2017. Development of a wearable device for motion capturing based on magnetic and inertial measurement units. *Scientific Programming* 2017 (2017).
- Oliver Glauser. 2019. Youtube video for Interactive Hand Pose Estimation using a Stretch-Sensing Soft Glove (SIGGRAPH 2019). <https://www.youtube.com/watch?v=Vrk4YmBRhac>
- Oliver Glauser, Shihao Wu, Daniele Panozzo, Otmar Hilliges, and Olga Sorkine-Hornung. 2019. Interactive hand pose estimation using a stretch-sensing soft glove. *ACM Transactions on Graphics (ToG)* 38, 4 (2019), 1–15.
- Frank L Hammond, Yiğit Mengüç, and Robert J Wood. 2014. Toward a modular soft sensor-embedded glove for human hand motion and tactile pressure measurement. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4000–4007.
- Shangchen Han, Beibei Liu, Randi Cabezas, Christopher D Twigg, Peizhao Zhang, Jeff Petkau, Tsz-Ho Yu, Chun-Jung Tai, Muzaffer Akbay, Zheng Wang, et al. 2020. MEgATrack: monochrome egocentric articulated hand-tracking for virtual reality. *ACM Transactions on Graphics (ToG)* 39, 4 (2020), 87–1.
- Shangchen Han, Beibei Liu, Robert Wang, Yuting Ye, Christopher D Twigg, and Kenrick Kin. 2018. Online optical marker-based hand tracking with deep labels. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–10.
- Shangchen Han, Po-chen Wu, Yubo Zhang, Beibei Liu, Linguang Zhang, Zheng Wang, Weiguang Si, Peizhao Zhang, Yujun Cai, Tomas Hodan, et al. 2022. UmeTrack: Unified multi-view end-to-end hand tracking for VR. In *SIGGRAPH Asia 2022 Conference Papers*. 1–9.
- Yihui He, Rui Yan, Katerina Fragkiadaki, and Shoubo Yu. 2020. Epipolar transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 7779–7788.
- Bingcheng Hu, Tian Ding, Yuxin Peng, Li Liu, and Xu Wen. 2020. Flexible and attachable inertial measurement unit (IMU)-based motion capture instrumentation for the characterization of hand kinematics: A pilot study. *Instrumentation Science & Technology* 49, 2 (2020), 125–145.
- 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.
- Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. 2018. Hand pose estimation via latent 2.5 d heatmap regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 118–134.
- Karim Isakov, Egor Burkov, Victor Lempitsky, and Yury Malkov. 2019. Learnable triangulation of human pose. In *Proceedings of the IEEE/CVF international conference on computer vision*. 7718–7727.
- Yifeng Jiang, Yuting Ye, Deepak Gopinath, Jungdam Won, Alexander W Winkler, and C Karen Liu. 2022. Transformer Inertial Poser: Real-time Human Motion Reconstruction from Sparse IMUs with Simultaneous Terrain Generation. In *SIGGRAPH Asia 2022 Conference Papers*. 1–9.
- David Kim, Otmar Hilliges, Shahram Izadi, Alex D Butler, Jiawen Chen, Iason Oikonomidis, and Patrick Olivier. 2012. Digits: freehand 3D interactions anywhere using a wrist-worn gloveless sensor. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. 167–176.
- Dong Uk Kim, Kwang In Kim, and Seungrul Baek. 2021. End-to-end detection and pose estimation of two interacting hands. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11189–11198.
- Hyuno Kim, Yuji Yamakawa, and Masatoshi Ishikawa. 2020. Robust hand tracking method by synchronized high-speed cameras with orthogonal geometry. In *2020 IEEE Sensors Applications Symposium (SAS)*. 1–5. <https://doi.org/10.1109/SAS48726.2020.9220075>
- Kyun Kyu Kim, Min Kim, Kyungrok Pyun, Jin Kim, Jinki Min, Seunghun Koh, Samuel E Root, Jaewon Kim, Bao-Nguyen T Nguyen, Yuya Nishio, et al. 2022. A substrate-less nanomesh receptor with meta-learning for rapid hand task recognition. *Nature Electronics* (2022), 1–12.
- Adarsh Kowdle, Christoph Rhemann, Sean Fanello, Andrea Tagliasacchi, Jonathan Taylor, Philip Davidson, Mingsong Dou, Kaiwen Guo, Cem Keskin, Sameh Khamis, David Kim, Danhang Tang, Vladimir Tankovich, Julien Valentin, and Shahram Izadi. 2018. The Need 4 Speed in Real-Time Dense Visual Tracking. *ACM Trans. Graph.* 37, 6, Article 220 (dec 2018), 14 pages. <https://doi.org/10.1145/3272127.3275062>
- Nikolaos Kyriazis and Antonis Argyros. 2014. Scalable 3d tracking of multiple interacting objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3430–3437.
- Dennis Laurijssen, Steven Truijen, Wim Saeys, and Jan Steckel. 2015. Three sources, three receivers, six degrees of freedom: An ultrasonic sensor for pose estimation & motion capture. In *2015 IEEE SENSORS*. IEEE, 1–4.
- Mengcheng Li, Liang An, Hongwen Zhang, Lianpeng Wu, Feng Chen, Tao Yu, and Yebin Liu. 2022. Interacting attention graph for single image two-hand reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2761–2770.
- Bor-Shing Lin, I-Jung Lee, Shu-Yu Yang, Yi-Chiang Lo, Junghsi Lee, and Jean-Lon Chen. 2018. Design of an inertial-sensor-based data glove for hand function evaluation. *Sensors* 18, 5 (2018), 1545.
- Kevin Lin, Lijuan Wang, and Zicheng Liu. 2021a. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1954–1963.
- Kevin Lin, Lijuan Wang, and Zicheng Liu. 2021b. Mesh graphormer. In *Proceedings of the IEEE/CVF international conference on computer vision*. 12939–12948.
- Yilin Liu, Shijia Zhang, and Mahanth Gowda. 2021. NeuroPose: 3D hand pose tracking using EMG wearables. In *Proceedings of the Web Conference 2021*. 1471–1482.
- Federico Lorussi, Enzo Pasquale Scilingo, Mario Tesconi, Alessandro Tognetti, and Danilo De Rossi. 2005. Strain sensing fabric for hand posture and gesture monitoring. *IEEE transactions on information technology in biomedicine* 9, 3 (2005), 372–381.
- Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. 2018. V2v-poseNet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5079–5088.
- Gyeongsik Moon, Shoubo Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. 2020. Interhand2.6M: A dataset and baseline for 3d interacting hand pose estimation from a single rgb image. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*. Springer, 548–564.
- Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. 2018. Generated hands for real-time 3d hand tracking from monocular rgb. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 49–59.
- Franziska Mueller, Micah Davis, Florian Bernard, Oleksandr Sotnychenko, Mickeal Verschoor, Miguel A Otaduy, Dan Casas, and Christian Theobalt. 2019. Real-time pose and shape reconstruction of two interacting hands with a single depth camera. *ACM Transactions on Graphics (ToG)* 38, 4 (2019), 1–13.
- Iason Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. 2011. Efficient model-based 3D tracking of hand articulations using Kinect. In *Bmvc*, Vol. 1. 3.
- Iasonas Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. 2012. Tracking the articulated motion of two strongly interacting hands. In *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 1862–1869.
- Timothy F O'Connor, Matthew E Fach, Rachel Miller, Samuel E Root, Patrick P Mercier, and Darren J Lipomi. 2017. The Language of Glove: Wireless gesture decoder with low-power and stretchable hybrid electronics. *PloS one* 12, 7 (2017), e0179766.
- Wookeun Park, Kyongkwan Ro, Suin Kim, and Joonbum Bae. 2017. A soft sensor-based three-dimensional (3-D) finger motion measurement system. *Sensors* 17, 2 (2017), 420.
- Yongbin Qi, Cheong Boon Soh, Erry Gunawan, and Kay-Soon Low. 2014. A wearable wireless ultrasonic sensor network for human arm motion tracking. In *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 5960–5963.
- Edoardo Remelli, Shangchen Han, Sina Honari, Pascal Fua, and Robert Wang. 2020. Lightweight multi-view 3d pose estimation through camera-disentangled representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 6040–6049.
- Javier Romero, Dimitrios Tzionas, and Michael J Black. 2022. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610* (2022).
- Yu Rong, Jingbo Wang, Ziwei Liu, and Chen Change Loy. 2021. Monocular 3D reconstruction of interacting hands via collision-aware factorized refinements. In *2021 International Conference on 3D Vision (3DV)*. IEEE, 432–441.
- Hochung Ryu, Sangki Park, Jong-Jin Park, and Jihyun Bae. 2018. A knitted glove sensing system with compression strain for finger movements. *Smart Materials and Structures* 27, 5 (2018), 055016.

- Tomohiko Sato, Shigeki Nakamura, Kotaro Terabayashi, Masanori Sugimoto, and Hiromichi Hashizume. 2011. Design and implementation of a robust and real-time ultrasonic motion-capture system. In *2011 International Conference on Indoor Positioning and Indoor Navigation*. IEEE, 1–6.
- Zhong Shen, Juan Yi, Xiaodong Li, Mark Hin Pei Lo, Michael ZQ Chen, Yong Hu, and Zheng Wang. 2016. A soft stretchable bending sensor and data glove applications. *Robotics and biomimetics* 3, 1 (2016), 22.
- Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. 2017. Hand keypoint detection in single images using multiview bootstrapping. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 1145–1153.
- Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. 2020. Weakly supervised 3d hand pose estimation via biomechanical constraints. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*. Springer, 211–228.
- Danhang Tang, Hyung Jin Chang, Alykhan Tejani, and Tae-Kyun Kim. 2014. Latent regression forest: Structured estimation of 3d articulated hand posture. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3786–3793.
- Danhang Tang, Jonathan Taylor, Pushmeet Kohli, Cem Keskin, Tae-Kyun Kim, and Jamie Shotton. 2015. Opening the black box: Hierarchical sampling optimization for estimating human hand pose. In *Proceedings of the IEEE international conference on computer vision*. 3325–3333.
- Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. 2016. Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision* 118 (2016), 172–193.
- Daniel Vlasic, Rolf Adelsberger, Giovanni Vannucci, John Barnwell, Markus Gross, Wojciech Matusik, and Jovan Popović. 2007. Practical motion capture in everyday surroundings. *ACM transactions on graphics (TOG)* 26, 3 (2007), 35–es.
- Timo Von Marcard, Bodo Rosenhahn, Michael J Black, and Gerard Pons-Moll. 2017. Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. In *Computer graphics forum*, Vol. 36. Wiley Online Library, 349–360.
- Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A Otaduy, Dan Casas, and Christian Theobalt. 2020. Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video. *ACM Transactions on Graphics (ToG)* 39, 6 (2020), 1–16.
- Robert Y Wang and Jovan Popović. 2009. Real-time hand-tracking with a color glove. *ACM transactions on graphics (TOG)* 28, 3 (2009), 1–8.
- Fu Xiong, Boshen Zhang, Yang Xiao, Zhiguo Cao, Taidong Yu, Joey Tianyi Zhou, and Junsong Yuan. 2019. A2j: Anchor-to-joint regression network for 3d articulated pose estimation from a single depth image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 793–802.
- Jan Xu, Weipeng Xu, Vladislav Golyanik, Marc Habermann, Lu Fang, and Christian Theobalt. 2020. EventCap: Monocular 3D Capture of High-Speed Human Motions using an Event Camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Xingchen Yang, Xueli Sun, Dalin Zhou, Yuefeng Li, and Honghai Liu. 2018. Towards wearable a-mode ultrasound sensing for real-time finger motion recognition. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 26, 6 (2018), 1199–1208.
- Xingchen Yang, Yu Zhou, and Honghai Liu. 2020. Wearable ultrasound-based decoding of simultaneous wrist/hand kinematics. *IEEE Transactions on Industrial Electronics* 68, 9 (2020), 8667–8675.
- Baowen Zhang, Yangang Wang, Xiaoming Deng, Yinda Zhang, Ping Tan, Cuixia Ma, and Hongan Wang. 2021. Interacting two-hand 3d pose and shape reconstruction from single color image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 11354–11363.
- Yang Zheng, Yu Peng, Gang Wang, Xinrong Liu, Xiaotong Dong, and Jue Wang. 2016. Development and evaluation of a sensor glove for hand function assessment and preliminary attempts at assessing hand coordination. *Measurement* 93 (2016), 1–12.
- Christian Zimmermann and Thomas Brox. 2017. Learning to estimate 3d hand pose from single rgb images. In *Proceedings of the IEEE international conference on computer vision*. 4903–4911.

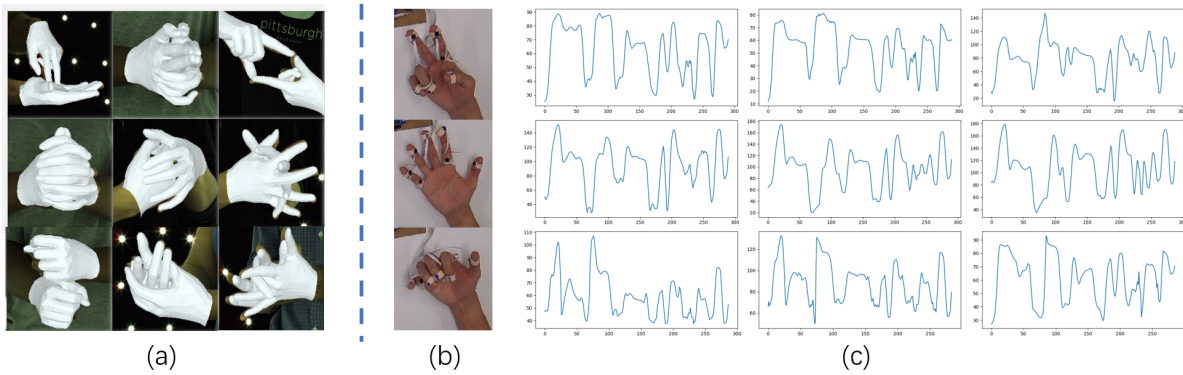


Figure 11: Dataset Visualization. (a) The InterHand 2.6M dataset used for pre-training. (b) Hand poses for fine-tuning. (c) Raw sensor data displaying variation as the hand moves to different poses. There are three subfigures in each line, representing the data feature dimensions 1-7, 8-14, and 15-21.

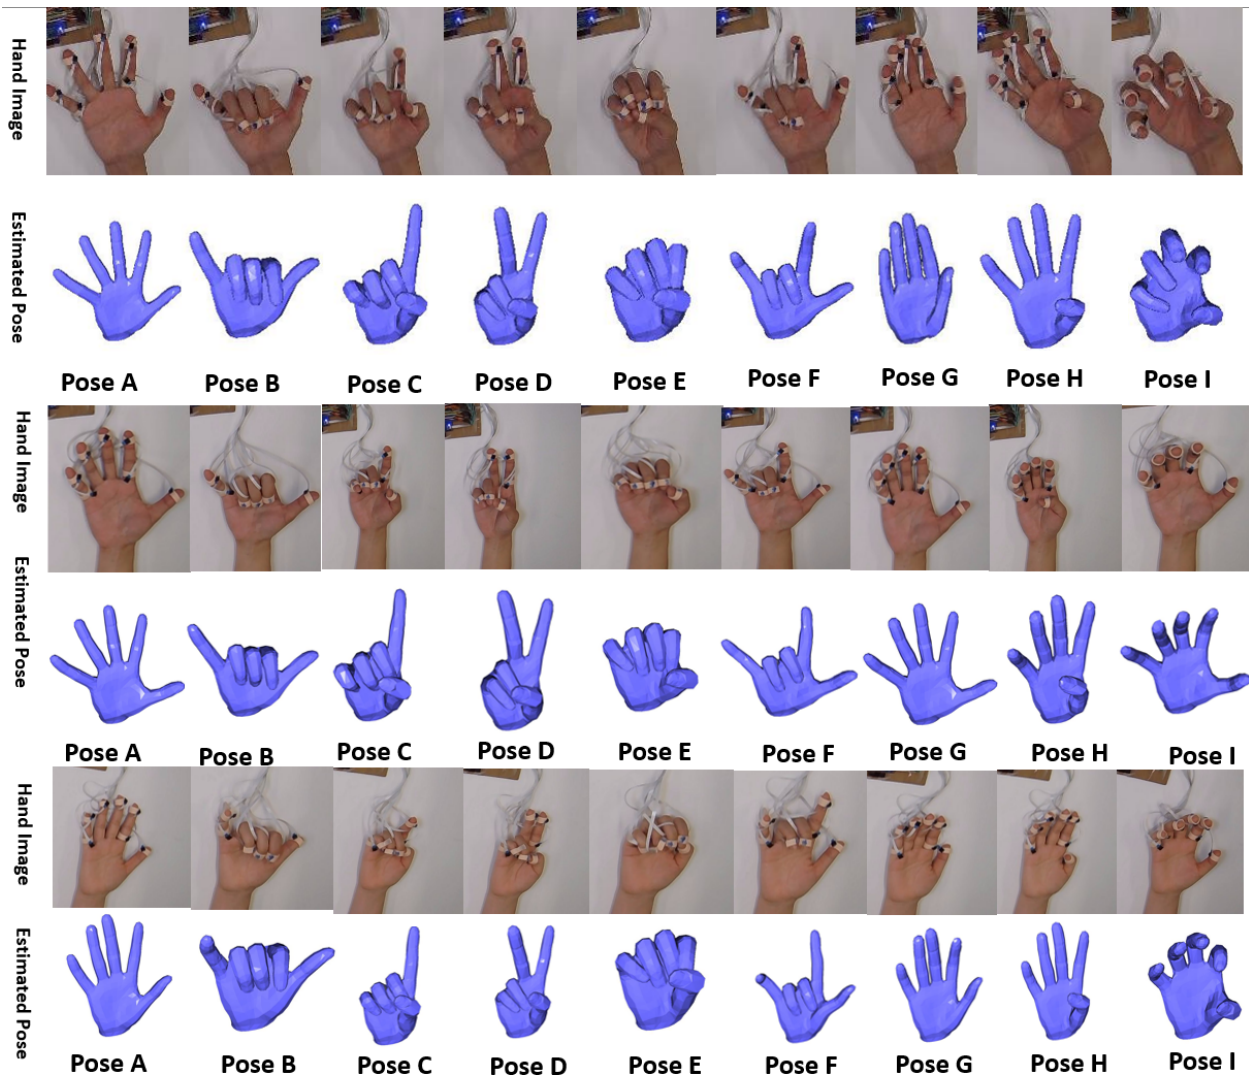
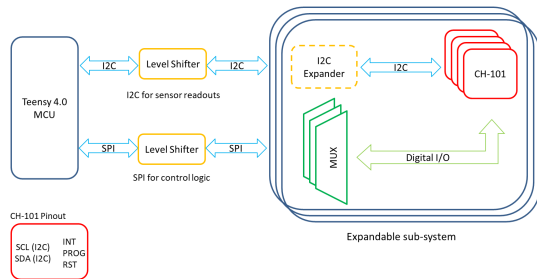


Figure 12: Evaluation of the sensitivity of results to hand size and shape. Here we show that the model fine-tuned on one dataset can be adapted to hands with different sizes and shapes. The top result is for the same hand size (“large”) as the data collected for fine-tuning, while the middle (“medium”) and lower (“small”) results show qualitatively similar performance for individuals with different hand sizes.

## SUPPLEMENTAL MATERIAL

### A TOWARDS FULL-DUPLEX OPERATION FOR HIGH REFRESH RATE

Real-time motion capture for precision movement reconstruction requires a high capture rate. With the advance of CMOS image sensors and high-speed photography, optical-based motion tracking has achieved a refresh rate in excess of 200 fps (frames per second) [Kim et al. 2020; Kowdle et al. 2018; Xu et al. 2020]. The ultrasound system demonstrated in this work has a refresh rate of around 10 Hz. The relatively low rate is due to delays in interrogating and waiting for replies from each sensor. To increase the refresh rate of the sensor front-end, FPGAs or dedicated controllers can be used to create multiple I2C buses for paralleled readback operation instead of sequential polling. Even faster operations can be achieved with a concept similar to that of FDMA. By having sensors sensitive to a wide spectrum of frequencies instead of one particular frequency, Therefore, all sensors can interrogate different frequencies simultaneously, allowing full-duplex operation and increase refresh rate significantly.



**Figure 13: System-level diagram for commercial MCU adaptation**

### B ADDITIONAL COMPARISONS WITH THE VISION BASELINE

We provide more analysis with the pure visual-based hand pose estimation model [Li et al. 2022], as shown in Fig 14. It is apparent

that the visual-based algorithm is sensitive to the background and the lighting conditions. When the background color is close to the human hand color, the model is collapsed.

### C TOWARDS SCALABLE SENSOR ACQUISITION SYSTEM ADAPTATION

Unlike the majority of commercially available sensors, CH-101 only supports 1.8V logic, while many commercial MCUs work at 3.3V or higher. Additionally, CH-101 uses a non-standard bi-directional drive mechanism that is not compatible with the popular push-pull or low-side open-drain counterparts. Special considerations are needed to address these problems.

As shown in Fig. 13, the proposed system contains multiple sub-units with each sub-unit supporting multiple CH-101 sensors. Each sub-unit consists of digital Muxes for logical I/O controls and an I2C expander for data readouts. Then each sub-unit's I2C and SPI buses are level shifted to 3.3V to match that of MCUs. Due to I2C buses' speed restriction (100kHz typical or 400kHz with high-speed mode), SPI-operated MUXs, which have a bus speed on the order of MHz, are chosen to handle control I/Os to avoid occupying the bandwidth. The optional I2C expander avoids the problem of multiple slaves having the same address and allows for another type of sensor to share the same I2C bus. The number of sensors each sub-unit supports depends on the MCU speed and I2C protocols. For higher-speed operations, an FPGA is preferred due to its reconfigurability.

Based on the system architecture mentioned above, we have conceptualized a schematic-level expandable system for general MCU integration (Fig. 15). The CH-101 consists of 3 additional I/O pins on top of I2C communication lines for resetting, triggering, and readback operations. For data telemetry, a dedicated I2C level shifter (TCA9406) is used. Pull-up resistors are placed on each side of the bus for open-drain operation. Correspondingly, several bidirectional level shifters (TXB0104) are used to level shift the SPI bus and other peripheral I/O lines for 1.8 environment. Three SPI-based I/O expanders (MCP23S08) handle the control of 8 CH-101s. "RST" and "PROG" pins are unidirectional, but "INT" pins are bidirectional with a non-industrial standard high-side open-drain drive, making them incompatible with the majority of the level shifters on the market. Therefore, additional buffers are added to convert the drive mechanism of the INT pin to the industrial-standard push-pull drives. Lastly, a power regulator (MIC5504) provides 1.8V power.

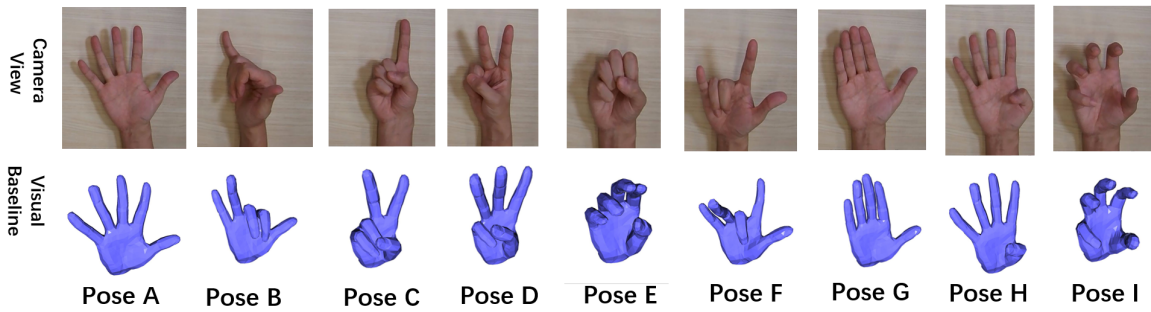


Figure 14: More visualization for visual-based baseline results. This visual-based algorithm is easily collapsed when the background color is close to the hand color. Meanwhile, the model is sensitive to the light conditions.

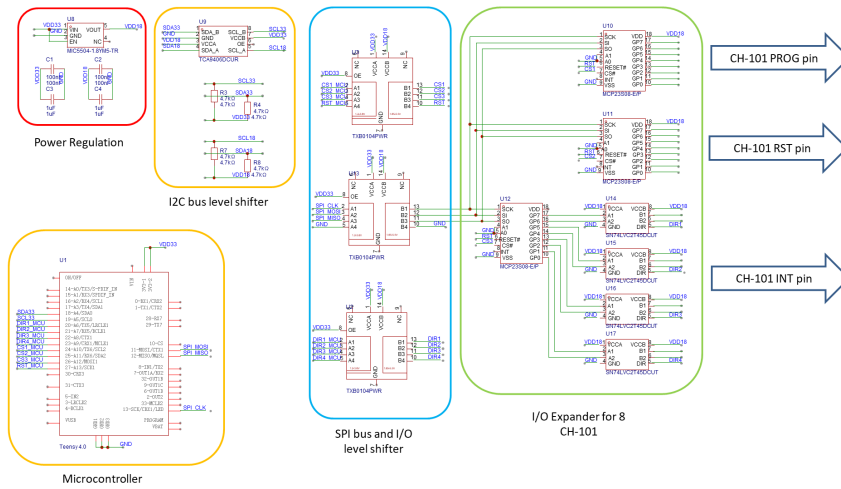


Figure 15: Our proposed circuit schematics for commercial MCU adaptation with scalable sensor arrays.