



# *mPose*: Environment- and subject-agnostic 3D skeleton posture reconstruction leveraging a single mmWave device

Cong Shi<sup>a</sup>, Li Lu<sup>b</sup>, Jian Liu<sup>c</sup>, Yan Wang<sup>d</sup>, Yingying Chen<sup>a,\*</sup>, Jiadi Yu<sup>e</sup>

<sup>a</sup> WINLAB, Department of Electrical and Computer Engineering, Rutgers University, New Brunswick, NJ, USA

<sup>b</sup> Key Laboratory of Blockchain and Cyberspace Governance of Zhejiang Province, Zhejiang University, Hangzhou, China

<sup>c</sup> Min H. Kao Department of Electrical Engineering and Computer Science, University of Tennessee, Knoxville, TN, USA

<sup>d</sup> Department of Computer and Information Sciences, Temple University, Philadelphia, PA, USA

<sup>e</sup> Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China

## ARTICLE INFO

### Keywords:

3D skeleton posture reconstruction  
mmWave

## ABSTRACT

Human skeleton posture reconstruction is an essential component for human–computer interactions (HCI) in various application domains. Traditional approaches usually rely on either cameras or on-body sensors, which induce privacy concerns or inconvenient practical setups. To address these practical concerns, this paper proposes a low-cost contactless skeleton posture reconstruction system, *mPose*, which can reconstruct a user's 3D skeleton postures using a single mmWave device. *mPose* does not require the user to wear any sensors and can enable a broad range of emerging mobile applications (e.g., VR gaming and pervasive user input) via mmWave-5G ready Internet of Things (IoT) devices. Particularly, the system extracts multi-dimensional spatial information from mmWave signals which characterizes the skeleton postures in a 3D space. To mitigate the impacts of environmental changes, *mPose* dynamically detects the user location and extracts spatial features from the mmWave signals reflected only from the user. Furthermore, we develop a deep regression method with a domain discriminator to learn a mapping between the spatial features and the joint coordinates of human body while removing subject-specific characteristics, realizing robust posture reconstruction across users. Extensive experiments, involving 17 representative body postures, 7 subjects, and 3 indoor environments, show that *mPose* outperforms contemporary state-of-the-art RF-based solutions with a lower average joint error of only ~30 mm, while achieving transferability across environments and subjects at the same time.

## 1. Introduction

Recent years have witnessed an upsurge of research in human movement tracking and recognition systems. These systems significantly facilitate human–computer interactions (HCI) and many emerging applications, such as virtual reality (VR), augmented reality (AR), fitness tracking, smart healthcare, and smart home control. Traditional approaches utilize vision-based techniques (Boulay, Bremond, & Thonnat, 2003; Mehrizi et al., 2018; Ren, Yuan, Meng, & Zhang, 2013) or body attached sensors (Chen & Xue, 2015; Jiang & Yin, 2015; VICON, 2019) to reconstruct skeleton postures and infer body movements. However, these approaches may suffer from illumination interference, incur privacy concerns, or introduce intrusive user experience. Therefore, a robust, privacy-preserving, and non-intrusive skeleton posture reconstruction method is highly demanded.

\* Corresponding author.

E-mail addresses: [cs1421@scarletmail.rutgers.edu](mailto:cs1421@scarletmail.rutgers.edu) (C. Shi), [li.lu@zju.edu.cn](mailto:li.lu@zju.edu.cn) (L. Lu), [jliu@utk.edu](mailto:jliu@utk.edu) (J. Liu), [y.wang@temple.edu](mailto:y.wang@temple.edu) (Y. Wang), [yingche@scarletmail.rutgers.edu](mailto:yingche@scarletmail.rutgers.edu) (Y. Chen), [jiadiyu@sjtu.edu.cn](mailto:jiadiyu@sjtu.edu.cn) (J. Yu).

<https://doi.org/10.1016/j.smhl.2021.100228>

Received 25 October 2021; Accepted 1 November 2021

Available online 25 November 2021

2352-6483/© 2021 Elsevier Inc. All rights reserved.

Non-intrusive approaches using radio frequency (RF) signals have been exploited to detect coarse-grained activities (Wang et al., 2014), monitor vital signs (Liu et al., 2015), and reconstruct body postures (Adib, Hsu, Mao, Katabi, & Durand, 2015; Jiang et al., 2020; Zhao, Li, et al., 2018). Among these studies, E-eyes (Wang et al., 2014) is a pioneer work to recognize indoor activities leveraging commodity WiFi devices (e.g., WiFi access points), while Liu et al. (Liu et al., 2015) first propose leveraging commodity WiFi to monitor vital signals including breathing rate and heart rate. WiTrack (Adib et al., 2015) can provide coarse tracking of body parts using customized antennas; RF-Pose (Zhao, Li, et al., 2018) has been developed to estimate 2D human poses using multiple universal software radio peripheral (USRP) units mounted on the walls. However, these systems require multiple dedicated devices (e.g., customized antennas, USRP devices), which are prohibitively expensive and hard to deploy in practice. WiPose (Jiang et al., 2020) demonstrates the feasibility of reconstructing 3D human posture using WiFi signals and deep learning, but it requires to deploy multiple antennas around the user, which has limited application scenarios.

As approaching the 5G era, low-cost Commercial Off-The-Shelf (COTS) mmWave hardware has been integrated into mobile devices (e.g., Soli on Google Pixel4 (Lien et al., 2016; *Don't hold the phone: new features coming to Pixel 4*, 2019)). Compared to WiFi, mmWave has a much shorter wavelength (i.e., about 1/16 of 5 GHz WiFi wavelength), which renders mm-level wireless ranging and enables high-precision skeleton reconstruction beyond traditional RF-based methods. Additionally, mmWave devices are usually equipped with miniature antenna arrays for highly directional transmissions, and such arrays can be used to derive the angles of arriving signals, which facilitates fine-grained spatial sensing. As such, researchers have utilized mmWave to recognize hand gestures (Lien et al., 2016), vital signs (Yang, Pathak, Zeng, Liran, & Mohapatra, 2017), and human identity (Yang, Liu, Chen, Guo, & Xie, 2020). Sengupta et al. (Sengupta, Jin, Zhang, & Cao, 2020) show the initial success of using mmWave radar to detect and track human skeletal postures. However, the proposed solution needs multiple mmWave devices and does not guarantee the generalizability to different environments (e.g., room layouts, furniture placement) and users (e.g., with different heights, lengths of arms and legs).

In this paper, we propose a low-cost contactless skeleton posture reconstruction system, *mPose*, which tracks a user's full-body 3D skeleton postures leveraging mmWave signals. Different from existing work, *mPose* can achieve fine-grained skeleton reconstruction with a single COTS mmWave device. By utilizing advanced signal processing and deep learning technologies, our system can dynamically remove posture-irrelevant information (e.g., reflections from static room objects and the user's location) and accurately track skeletal joints (e.g., wrists, elbows, knees, and ankles) in a 3D space across different domains (i.e., with different users or in different environments.) *mPose* can be conveniently deployed in a 5G or 802.11ad-enabled IoT device (e.g., a smartphone or a smart TV) without additional costs. It provides privacy-preserving tracking of fine-grained human poses, facilitating various applications, such as contactless smart appliance control, augmented reality and virtual reality games, fitness tracking, wellbeing monitoring, and smart city surveillance.

Realizing such a skeleton posture reconstruction system has a number of challenges in practice. First, it is challenging to reconstruct 3D full-body skeleton posture using a single device. Though mmWave signal can provide fine-grained information, existing approaches usually require at least two devices to achieve it (Sengupta et al., 2020). Second, skeleton reconstruction usually occurs at different locations in a room or in different rooms with distinct layouts and furniture placement. Such heterogeneous environmental factors may cause different reflections that could significantly impact the skeleton reconstruction performance. Third, it is essential to have a general skeleton reconstruction model so that individual users do not need to suffer tedious training processes. However, due to differences in individual body shapes, it is difficult to train such a model to achieve high skeleton reconstruction accuracy for different users.

To address these challenges, we propose to extract fine-grained spatial information from mmWave signals reflected by human bodies to perform spatial tracking of the users' body parts. Specifically, Frequency-Modulated Continuous Wave (FMCW) (Stove, 1992) is used to obtain the range of the user's body parts. Capon beamforming is used with an antenna array to derive angles of arriving signals, which are integrated with the range information to enable spatial tracking of human body parts.

To enhance the resolution of spatial tracking, we exploit the virtual antenna technology (Rao, 2017) to develop a virtual antenna array with a much larger size than the physical antenna array. In addition, we dynamically estimate the user's location and extract the spatial features from the signals reflected from the user in a 3D contour, which mitigates the impacts from different environmental setups and user locations. A three-layer convolutional neural network (CNN) is designed to model human skeleton postures based on the extracted spatial features. Furthermore, *mPose* employs domain adversarial learning (Ganin et al., 2016) to remove user-specific characteristics embedded in the spatial features and ensure reliable skeleton reconstruction across users. The major contributions of this work are summarized as follows:

- We show that it is possible to accurately reconstruct 3D full-body skeleton postures by using a single COTS mmWave device. We develop a low-cost 3D skeleton posture reconstruction system, *mPose*, which can precisely localize skeletal joints in a 3D space.
- We develop a target detection method that allows our system to extract spatial features from the signals reflected only by the user and mitigate the impacts from the changes of environments and user locations.
- We design a domain discriminator that removes user-specific characteristics embedded in the mmWave signals so as to achieve robust skeleton reconstruction across users with reasonable training efforts.
- We conduct extensive experiments involving 17 body postures, 7 subjects, and 3 different environments to evaluate the performance of *mPose*. The results demonstrate that our system can achieve a low average joint estimation error of 30 mm while achieving environment- and user-independency.

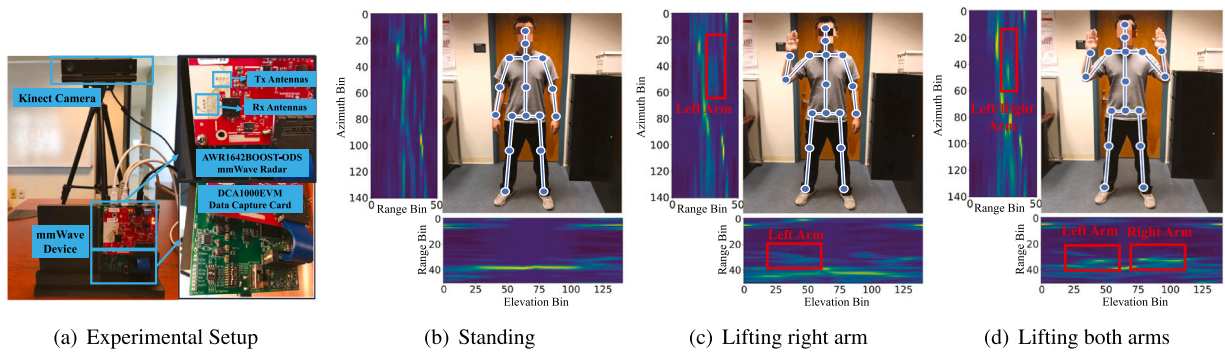


Fig. 1. mmWave device and the received signals impacted by different postures.

## 2. Related work

Existing skeleton posture reconstruction studies usually rely on either attached sensors or cameras (Cao, Hidalgo, Simon, Wei, & Sheikh, 2018; Ho et al., 2016; Iwasawa et al., 2000; Mehrizi et al., 2018; Mueller et al., 2019; Quwaider & Biswas, 2008; Schwarz, Mkhitarian, Mateus, & Navab, 2012; Shen, Wang, & Roy Choudhury, 2016; Shen et al., 2016; Thompson, Jain, LeCun, & Bregler, 2014; Valero, Sivanathan, Bosché, & Abdel-Wahab, 2017; VICON, 2019). For instance, Quwaider et al. (Quwaider & Biswas, 2008) exploit inertial sensors attached to the human body and Hidden Markov Model (HMM) to reconstruct body postures. Shen et al. (Shen et al., 2016) develop a system that can track 3D postures of an entire arm (i.e., both wrist and elbow) using motion and magnetic sensors on smartwatches. While these approaches are intuitive, they all require users to wear additional sensors, which is not always feasible in practice. There are also a number of studies that utilize cameras for 2D/3D skeleton posture reconstruction (Mehrizi et al., 2018; Thompson et al., 2014). For example, Thompson et al. (2014) design a spatial-model based on a deep convolutional neural network to map human postures in a video frame into 2D joint positions. By utilizing a camera array, Mehrizi et al. (2018) demonstrate the potential to recover 3D human postures through analyzing image features of joint kinematics. However, these vision-based approaches usually incur privacy concerns, and they require line-of-sight and sufficient illumination, which largely limits their application scenarios.

To overcome the weaknesses of body-attached sensor- and vision-based solutions, researchers have been exploring contact-less sensing approaches to recognize human activities (Postolache, Girão, Madeira, & Postolache, 2010; Wang et al., 2014; Xu et al., 2017), monitor vital signs (Liu et al., 2015), or reconstruct human body postures (Adib et al., 2015; Feng et al., 2020; Jiang et al., 2020; Papandreou, Zhu, Kanazawa, Toshev, Thompson, Bregler, & Murphy, 2017; Wang, Zhou, Panev, Han, & Huang, 2019; Zhao, Li, et al., 2018). For example, E-eyes (Wang et al., 2014) is one of the first studies to recognize various daily activities (e.g., cooking, eating, and watching TV) using commodity WiFi devices. Liu et al. (2015) developed the first non-intrusive vital sign tracking system that can monitor users' breathing and heart rates during sleep using WiFi signals. To reconstruct body postures, Zhao, Li, et al. (2018) develop a system that uses customized antenna arrays and FMCW signals to estimate 2D postures. In addition, Adib et al. (2015) propose to track the human arm motions leveraging FMCW radio operating on a universal software radio peripheral (USRP) device. However, these systems require dedicated RF settings (e.g., USRP devices and antenna arrays), which are costly and require to be installed in controlled environments. A recent work shows that channel state information (CSI) of WiFi signals could be used to reconstruct 3D human postures (Jiang et al., 2020), but the system requires a large number of antennas distributed around the human body, making it hard to deploy in real-world scenarios. Additionally, due to the open-propagation nature, WiFi signals are very sensitive to various environmental factors (e.g., the room layout and moving objects), which inevitably cause reconstruction errors.

Compared to existing work, *mPose* leverages highly directional mmWave signals to track skeletal joints and thus allows more reliable posture construction. The shorter wavelengths of mmWave signals also enable posture tracking with higher precision. A more recent study (Sengupta et al., 2020) shows the initial success of using mmWave signals to track human skeleton postures, but it requires two distributed mmWave devices to achieve 3D posture tracking and the performance of the system is moderate. In contrast, *mPose* could achieve high accuracy in reconstructing 3D postures by using a single mmWave transceiver. In addition, our system dynamically extracts spatial features from the target user to remove the impacts of user location and can be adapted to reconstruct postures across users with a reasonable amount of training efforts.

## 3. Preliminaries

In this section, we first introduce the fundamentals of mmWave sensing. Then, we investigate the feasibility of using mmWave signals for skeleton posture reconstruction. We also present some potential applications that can be benefited from *mPose*.

### 3.1. Fundamentals of mmWave sensing

Recently, mmWave-based sensing techniques have been widely investigated in academia and industry. Due to the mmWave signals' high frequency (i.e.,  $> 30$  GHz) and short wavelength (i.e.,  $< 10$  mm), mmWave techniques can realize high-resolution measurements for fine-grained sensing. Moreover, mmWave transceivers can be implemented in a compact form that is small enough to fit into mobile devices (e.g., Google Pixel series (*Don't hold the phone: new features coming to Pixel 4*, 2019)). As such, we propose to employ a single COTS mmWave device (i.e., TI AWR1642BOOST-ODS mmWave device), which is low-cost, portable, and highly-integrated, for 3D human posture reconstruction in this work. The mmWave device supports multiple-input and multiple-output (MIMO) with an antenna array (e.g.,  $2 \times 4$  antennas as shown in Fig. 1(a)) that can provide multiple dimensions of temporal and spatial information with one single measurement.

To reconstruct a user's skeleton posture, *mPose* needs to estimate the range (distance) and angle of different human body parts to the mmWave device. Particularly, our system utilizes the mmWave device to transmit a continuous chirp signal sweeping across a bandwidth within a fixed duration. Upon receiving the reflected mmWave signals, the mmWave device performs a dechirp operation on both the transmitted and the received signals to derive an intermediate frequency (IF) signal, which can be formulated as:

$$IF = \sin[(f_{Tx} - f_{Rx})t + (\phi_{Tx} - \phi_{Rx})], \quad (1)$$

where  $f_{Tx}$  and  $f_{Rx}$  represent the instantaneous frequencies of the transmitted and received mmWave signals, respectively. While  $\phi_{Tx}$  and  $\phi_{Rx}$  denote the instantaneous phases of the two signals respectively. By integrating IF signals from multiple antennas of the mmWave device, *mPose* can extract spatial information (i.e., range and angle) on a user's skeletal posture. We introduce how to use the IF signal to derive spatial features in Section 5.

### 3.2. Feasibility of reconstructing skeleton postures via a single mmWave device

To demonstrate the feasibility of utilizing mmWave signals to reconstruct 3D human postures, we conduct a preliminary experiment by examining the spatial information derived from a single portable mmWave device when a volunteer is performing a set of postures (i.e., standing, waving right arm, and waving both arms). Specifically, we use TI AWR1642BOOST-ODS shown in Fig. 1(b) as the mmWave transceiver. A Microsoft Kinect (Zhang, 2012) adjacent to the mmWave device is used to record the video as the ground truth 3D skeleton joint positions. Fig. 1(b) (c), and (d) show the captured video frames and the corresponding range-azimuth and range-elevation heatmaps of the three postures. The heatmap value represents the frequency response of mmWave at different angles (i.e., azimuths and elevations) and ranges (i.e., distances), with higher values representing stronger reflections. Across the three postures, we can find the contour of torso and limb across the range bin  $25 \sim 40$  (i.e.,  $1.1 \text{ m} \sim 1.8 \text{ m}$ ), which is consistent with our ground truth. Furthermore, as the volunteer lifts the right arm as shown in Fig. 1(c), intensity changes at the range bin 25 and azimuth bin 30 can be observed. Interestingly, in Fig. 1(d), we can observe symmetric intensity changes in the range-elevation heatmap, representing the posture of lifting both arms. The preliminary study shows the feasibility of using the mmWave signals from a single device to capture the spatial information of different human body parts.

### 3.3. Potential applications

*mPose* can reveal human full body skeleton postures in a fine-grained manner, which can be utilized to support an abroad array of applications.

**Alerting acts of violence based on human posture:** some existing studies demonstrated the potential of interpreting human activities based on the posture (Guo, Yu, Ding, & Skitmore, 2018; Ho et al., 2016). As public safety raises increasing concern recently, inferring the activities of violence (e.g., sexual violence, physical fighting) by embedding *mPose* into public infrastructures (e.g., street lamps) is highly desirable. Upon detecting the violence, *mPose* can alert the perpetrator and contact the police to prevent injury, psychological harm, or even death. Different from vision-based approaches, our system will not disclose the privacy of human subjects since the personally identifiable information (e.g., facial information) is not embedded in mmWave signals. Furthermore, our mmWave-based system will not suffer from performance degradation due to complicated ambient lighting conditions.

**Full-body AR/VR gaming:** the unique immersive experience provided by AR/VR gaming has attracted millions of users around the world. The AR/VR games create a 3D virtual world and allow users to use hand-held controllers to interact with 3D objects in the virtual world (Caserman, Garcia-Agundez, Konrad, Göbel, & Steinmetz, 2019; HTV, 2021). To improve user experience, *mPose* can facilitate AR/VR gaming by enabling a controller-free paradigm. The miniature mmWave sensors can be easily embedded into the AR/VR devices, such as AR/VR headsets, game consoles, or even smartphones.

## 4. Overview of mPose

### 4.1. Challenges

**3D Skeleton Posture Reconstruction Using a Single mmWave Device.** Compared to the existing solutions that use multiple devices or spatial-distributed antenna arrays, our single-device approach is challenging as we only have one device equipped with an on-board antenna array for sensing, which provides limited spatial diversity and a lower sensing resolution. Thus, it is crucial to extract in-depth features that can capture the unique impacts of different postures from mmWave signals.




Fig. 2. System flow of mPose.

**Location and Environment Changes.** In real-world scenarios, people may perform postures at different locations of a room or even in different rooms. Although mmWave signals reflected from the human body can be used to infer their relative distance and angle to the mmWave device, such estimation is prone to the varying location (e.g., causing different distances and angles between the subject and the mmWave device). Furthermore, the mmWave signals reflected from the objects in the room (e.g., furniture and home appliances) or walls can introduce interference to the received mmWave signals, leading to joint position estimation error. Our system needs to derive robust feature representations for accurate posture reconstruction.

**Body Characteristic Differences Across Subjects.** The mmWave signals also carry substantial information specific to individual subjects (e.g., body characteristics such as heights, length of arm and leg). Such body characteristics variations make it difficult to train a general skeleton reconstruction model applicable to different subjects, since the mapping between the spatial features and the joint positions will not be consistent. To reduce training efforts during practical deployment, it is highly desired to enable accurate posture reconstruction across users.

#### 4.2. System architecture

In this work, we address the above challenges and develop a 3D skeleton posture reconstruction system, *mPose*, which can continuously track a user's joints using a single COTS mmWave device. *mPose* is designed to examine the mmWave signals to determine the spatial positions of joints for posture reconstruction. The system flow is illustrated in Fig. 2. Specifically, in our system, the *Data Preprocessing & Feature Extraction* module first demodulates the reflected mmWave signals to obtain the ranges of various reflectors (e.g., the user's body parts, room objects), and then integrates the range information across antennas on the mmWave device to derive the angle information (i.e., azimuth and elevation). A virtual antenna array is constructed to enhance the resolution of the derived angle information, which enables fine-grained skeleton posture tracking on a single mmWave device. To ensure the robustness to user location and environment changes, the system dynamically tracks the user's location and extracts spatial features from the mmWave signals reflected from the user, which removes the impacts of the user's locations (i.e., the relative distances and angles between the user and the mmWave device) and environmental factors (e.g., reflections from walls and room objects).

After obtaining the spatial features, we develop a *Subject-independent Skeleton Posture Reconstruction Model* to map the spatial features to 3D coordinates of a user's skeletal joints (e.g., wrists, elbows, knees, and ankles) while removing user-specific body characteristics embedded in the spatial features. Instead of directly inferring the joint positions based on the spatial features, we resort to a deep-learning-based approach to learn a more reliable mapping between the joint positions and the spatial features.






Fig. 3. Illustration of the antenna arrays used in *mPose*.

Particularly, we use the joint positions captured with Kinect (Zhang, 2012) to train a deep regression model, which is referred to as a 3D posture model. The 3D posture model consists of a representation extractor that derives skeleton representations from spatial features and a 3D posture estimator that predicts 3D coordinates of 17 skeletal joints. In addition, *mPose* also employs a domain discriminator to remove subject-specific characteristics in the spatial features via domain adversarial training and adjust the 3D posture model to enable general skeleton reconstruction across users.

## 5. Design of *mPose*

### 5.1. Mmwave signal design

We need to design the FMCW signal (i.e., a chirp signal) in a way that maximums the sensing power of *mPose*. In other words, we want to simultaneously achieve high range and angle resolutions for fine-grained skeleton reconstruction. The range resolution  $d_{res}$  of the FMCW signal is defined as (Iovescu & Rao, 2017):

$$d_{res} = \frac{c}{2 \times B}, \quad (2)$$

where  $c$  is the speed of light;  $B$  is the bandwidth of a mmWave chirp signal. According to the formulation, the range estimation resolution is determined by the bandwidth (i.e.,  $B$ ). Hence, the bandwidth of the mmWave chirp signal should be set as large as possible. Since the frequency band used by the mmWave device lies in the range of 77 ~ 81 GHz, the chirp bandwidth of *mPose* is selected as 81–77=4 GHz, with a range estimation resolution of 3.75 cm. Additionally, we use a short chirp duration  $\tau = 33 \mu s$  to maintain the capability of *mPose* on continuous sensing. The angle resolution of mmWave signals can be formulated as:

$$\theta_{res} = \frac{\lambda}{N \times d}, \quad (3)$$

where  $\lambda$  denotes the wavelength;  $N$  and  $d$  are the number of receiving antennas and the distance between two neighboring antennas in the antenna array, respectively. Since the mmWave device used in *mPose* is equipped an antenna array with  $N = 4$  and  $d = \frac{\lambda}{2}$ , the angle resolution is  $0.5^\circ$ . To enable fine-grained posture reconstruction, we synthesize a virtual antenna array to double the angle resolution, with the details elaborated in Section 5.2.

### 5.2. 3D Spatial feature extraction

To reconstruct the skeleton posture from mmWave signals, *mPose* first estimates the range (i.e., distance) between each of the human body parts and the mmWave device. The principle of extracting such range information is to measure the frequencies of IF signals, which are linearly correlated with the distances to reflectors as mentioned in Section 3.1. Specifically, we apply range-FFT on the IF signal to obtain the frequencies associated with strong reflections (i.e.,  $f_{IF}$ ), such as those from the subject's body parts and room objects, and then map the frequencies to distances:

$$d = \frac{f_{IF} \times c \times \tau}{2B}, \quad (4)$$

where  $c$  denotes the speed of light;  $B$  and  $\tau$  are the bandwidth and the propagation time of the chirp signal, respectively. To reconstruct the user's posture in a 3D space, we further estimate the phase shifts (i.e.,  $\omega$ ) associated with the Angle-of-Arrival (AoA) of the reflectors. Particularly, we apply angle-FFT upon the range information derived from all receiving antennas to obtain  $\omega$ , which can be mapped to angles (e.g., azimuths or elevations):

$$\theta = \sin^{-1}\left(\frac{\lambda \times \omega}{2\pi \times \Delta d}\right), \quad (5)$$

where  $\Delta d$  denotes the distance between two receiving antennas and  $\lambda$  represents the wavelength. The angle and range information together serve as the basis of posture reconstruction in *mPose*.

To enhance the resolution of angle estimation, *mPose* exploits a  $1 \times 8$  antenna array with the diagram shown in Fig. 3. Besides 4 physical antennas (i.e., yellow circles), *mPose* also simulates 4 virtual antennas (i.e., green circles), which doubles the angle resolution of the antenna array. Particularly, our system pairs  $T \times 1$  with  $R \times 1 \sim R \times 4$  to form a  $1 \times 4$  physical array. In addition, by pairing the other physical transmitting antenna (i.e.,  $T \times 2$ ) with the same 4 physical receiving antennas, we can synthesize a  $1 \times 4$  virtual array. The virtual array can operate together with the physical one by using time-division multiplexing, which alternatively transmits mmWave chirps signals using the physical transmitting antennas (i.e.,  $T \times 1$  and  $T \times 2$ ) and receives the signals using the 4 physical receiving antennas (i.e.,  $R \times 1 \sim R \times 4$ ). Given the 8 receiving antennas, our system can achieve an angle resolution of  $0.25^\circ$ .

Furthermore, to track skeletal joints in a Cartesian coordinate system, the antennas are organized into orthogonal subarrays, with two subarrays placed horizontally to capture range-azimuth and one subarray put vertically to derive range-elevation. Fig. 3 illustrates such a spatial relationship (i.e., range, azimuth, and elevation) between the mmWave device and the torso of a subject (i.e., performing a standing posture). By sequentially performing range-FFT and angle-FFT on mmWave signals received with an antenna array, we derive a 2D heatmap that captures the frequency response of signals reflected from different angles at different distances:

$$F(d, \theta) = \text{AngleFFT}(\text{RangeFFT}(IF, d), \theta), \quad (6)$$

where  $\text{RangeFFT}(\cdot, d)$  and  $\text{AngleFFT}(\cdot, \theta)$  represent the FFT operations at range  $d$  and angle  $\theta$ , respectively. We then combine multiple such 2D heat-maps (i.e., 2 range-azimuth heatmaps and 1 range-elevation heatmap) to derive 3D spatial features for posture reconstruction.

### 5.3. Environmental interference removal via target detection

Since the mmWave signals capture the spatial information of all reflectors within the field of view, the static objects in a room (e.g., furniture, walls) can introduce a significant amount of posture-irrelevant noises which interfere with our posture reconstruction system. In addition, the variations of the subject's location can alter the relative distance and angle between the subject and the mmWave device, introducing uncertainty to the derived spatial information of user posture. To mitigate these two types of interferences, *mPose* dynamically detects the 3D location of the target user's torso and creates a 3D contour to segment the spatial information (i.e., range-azimuth and range-elevation heatmaps), which can be used to remove the impacts of location and environment factors.

Particularly, *mPose* tracks the torso's 3D location through examining the reflection energy in the range and angle bins, since the torso normally leads to the strongest energy due to its larger reflection area compared to arms/legs. We denote the range, azimuth, and elevation of the torso as  $(r, \theta, \psi)$ . After obtaining the torso's location, *mPose* removes the location impacts by extracting the spatial information across ranges:  $[r - \Delta r, r + \Delta r]$ , where  $\Delta r$  is determined by the maximum arm length of the user. To remove the environmental impacts, we dynamically calculate the angles where the human body resides based on  $r$  to further segment the spatial information. Specifically, we extract the heatmaps across azimuth bins:  $[\theta - \Delta\theta, \theta + \Delta\theta]$ . The variable  $\Delta\theta$  is calculated through:

$$\Delta\theta = \arctan^{-1} \frac{l_{arm}}{2r}, \quad (7)$$

where  $l_{arm}$  is the user's arm span. Similarly, we calculate the elevation variable  $\Delta\psi$  based on the user's height and extract the spatial information across the derived elevation bins. Since the varying distance between the subject and the mmWave device could change the size of the extracted range-azimuth and range-elevation heatmaps, we interpolate the heatmaps to a fixed dimension to mitigate the impacts of such variations, which ensures consistent heatmap sizes under the distance variations. With the approach, we can eliminate the interferences from the location and the environment in the heatmaps. Fig. 4 shows a heatmap example before and after applying the proposed approach. We can find that compared with Fig. 4(a), the noises are significantly mitigated in Fig. 4(b), which confirms the effectiveness of our approach.

### 5.4. Subject-independent 3D skeleton posture reconstruction

**Model Overview.** In order to construct 3D human posture, we build a deep learning model to map the extracted range-azimuth and range-elevation heatmaps to the coordinates of joints in the 3D space. Fig. 5 shows the architecture of the proposed deep learning model. The deep learning model takes data  $X$  (i.e., extracted spatial information) and the 3D posture ground truth  $P$  derived by Microsoft Kinect as input. The ground truth includes the 3D positions of 17 skeletal joints (e.g., wrists, elbows, knees) (Zhang, 2012). The input data are first converted into a set of low-rank posture representations  $Z$  using the posture representation extractor implemented by a 3-layer CNN model. Based on the posture representations, a 3D posture estimator can predict the 3D joint coordinates (i.e.,  $\hat{P}$ ). In addition, a domain discriminator that can predict the subject label  $D$  is used to assist in training the posture representation extractor and posture estimator. By optimizing the posture representation extractor with the domain discriminator to achieve indistinguishable subject labels (i.e., maximizing the domain loss), our model can remove the subject-specific characteristics for general skeleton reconstruction across users.

**3D Posture Reconstruction Model.** The posture representation extractor is a 3-layer CNN model. Particularly, we use a convolutional layer with 2D filters in each CNN layer to calculate the feature maps. By integrating the 2D features maps of range-azimuth and range-elevation information, we can derive posture representations (i.e.,  $Z$ ) that characterize skeleton posture in the 3D space. In addition, a max-pooling layer is attached to the convolutional layer. Max-pooling can enhance the representation




Fig. 4. Illustration of the environmental interference removal through detecting the target user in a 3D contour.




Fig. 5. Deep learning architecture used in *mPose* for subject-independent posture reconstruction.

transferability by integrating multiple 2D feature points over a small neighborhood. To prevent over-fitting, we use a dropout layer to randomly remove network parameters during training. The 2D feature maps are then flattened and compressed with a fully-connected layer. Given input data  $X$ , the posture representation extractor produces 3D posture representations as follows:

$$Z = F(X, \Phi), \quad (8)$$

where  $F(\cdot)$  denotes the CNN model and  $\Phi$  represents the trainable parameters (i.e., model weights) of the posture representation extractor. The model implementation of *mPose* employs 64, 128, and 256 filters in the 3 convolutional layers, respectively. The dropout rate of the dropout layer is set to 50%. We use Leaky ReLU with the parameter  $\alpha = 0.01$  for both the convolutional and the fully-connected layers.

Based on the derived representations  $Z$ , *mPose* employs a 2-layer fully-connected neural network (i.e., posture estimator) to estimate the 3D position of each skeletal joint. The neural network further extracts non-linear abstractions that characterize the 3D skeleton posture. Based on the abstractions, a regression layer produces the joint coordinates  $\hat{P}$ . We define the mapping function as:

$$\hat{P} = G_p(Z, \Theta), \quad (9)$$

where  $G_p(\cdot)$  represents the neural network and  $\Theta$  denotes its trainable parameters. To optimize the posture estimator for predicting the 3D skeletal joint positions, we use Huber loss (Wikipedia, 2019) to quantify the estimation cost. The joint estimation loss is defined as:

$$L_p = \begin{cases} \frac{1}{2}(P - \hat{P})^2, & |P - \hat{P}| \leq \delta \\ \delta \cdot (|P - \hat{P}| - \frac{1}{2}\delta), & |P - \hat{P}| > \delta \end{cases} \quad (10)$$

where  $P$  is the ground truth of the joint coordinates, and  $\delta$  is the threshold for outlier detection.

**Domain Discriminator.** To achieve robust posture reconstruction across subjects, we use domain adversarial training (Ganin et al., 2016) to optimize the posture representation extractor. The core component is a domain discriminator used in the training process for removing the domain-specific characteristics entangled in posture representations. Specifically, the domain discriminator is a 2-layer fully-connected neural network taking the posture representations  $Z$  as input. It infers the subject label as:

$$\hat{D} = G_d(Z, \Gamma), \quad (11)$$






Fig. 6. Illustration of the three environments used for evaluation.

where  $G_d(\cdot)$  denotes the neural network and  $\Gamma$  represents the corresponding trainable parameters. We optimize the neural network with the cross-entropy loss as:

$$L_d = H(D, \hat{D}), \quad (12)$$

where  $H(\cdot)$  represents the cross-entropy loss function, and  $D$  denotes the ground truth of the domain label.

The domain discriminator seems to contradict with our objective of domain-independent posture reconstruction. However, by using an adversarial loss, we can optimize the posture representation extractor to fool the domain discriminator so as to make the derived representations domain-independent. Specifically, during training, we apply a negative factor  $-\lambda$  to the domain loss to force the posture representation extractor to maximize the domain loss. We define the adversarial loss to optimize the posture representation extractor as follows:

$$L_{adv} = L_p - \lambda L_d, \quad (13)$$

where  $L_p$  and  $L_d$  are the joint estimation loss and the domain loss, respectively. The factor  $\lambda$  is used to balance the performance of domain adaptation and posture reconstruction. In the adversarial training process, we iteratively optimize  $\{\Phi, \Theta\}$  and  $\Gamma$  with Adam optimizer.

## 6. Evaluation

### 6.1. Experimental setup & methodology

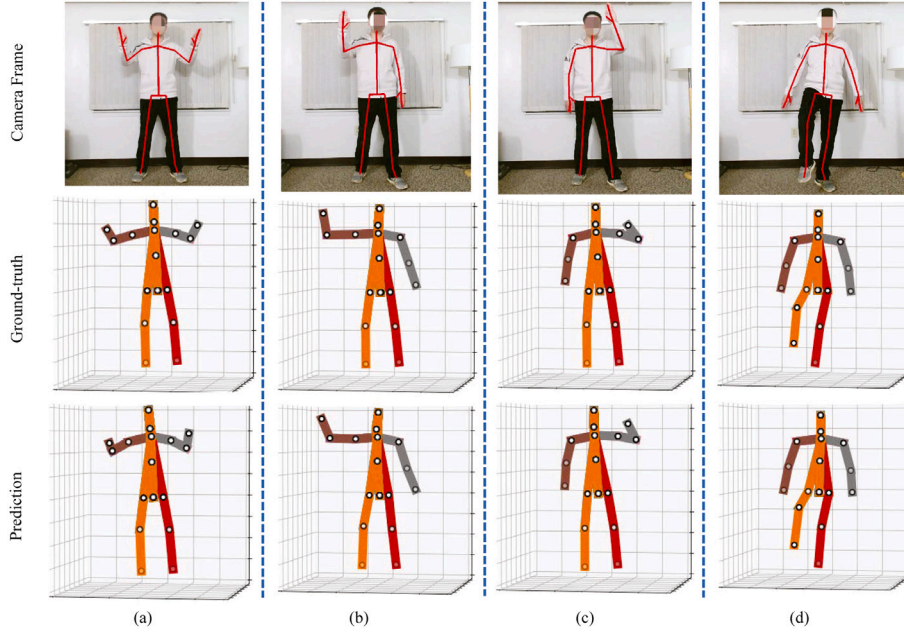
**Hardware Components.** We implement *mPose* with a TI AWR1642-ODS mmWave device as the sensing front end and a ThinkPad X1 Carbon laptop as the data processing backend. The mmWave device has an integrated antenna array with 2 transmitting antennas and 4 receiving antennas, which sends and receives chirp signals with a frequency range of 77 ~ 81 GHz and a chirp duration of 33  $\mu$ s. A TI DCA1000EVM data capture card is used to collect data from the mmWave device and forward them to the laptop. While we are collecting mmWave data, we use a Microsoft Kinect (Zhang, 2012) to record the ground truth 3D coordinates of 17 skeletal joints (i.e., the same set of joints evaluated in WiPose (Jiang et al., 2020)).

**Data Collection.** Experiments are conducted in three environments of different sizes and room objects as shown in Fig. 6 to demonstrate the robustness of *mPose* across environments. The larger room (i.e., lab) is a public place and has a size of 28 ft  $\times$  25 ft with desks, chairs, and many lab devices (e.g., desktops, 3D printers). The two smaller rooms (i.e., office and apartment) have sizes of 24 ft  $\times$  15 ft and 33 ft  $\times$  17 ft with office (e.g., tables, chairs) and home objects (e.g., sofas, floor lamps). We use the three rooms of different sizes and layouts to demonstrate the effectiveness of the proposed environmental interference removal algorithm. We recruited 7 subjects, including 5 males and 2 females with various heights and weights (e.g., the range of height and weight are 1.63 ~ 1.78 m and 58 ~ 83 kg respectively). The experiments involved 17 representative postures: lifting left/right arm to the front for 45/90/180 degree, lifting left/right arm from the side for 45/90/180 degree, lifting left/right leg for 45/90 degree, waving hands, walking, random moving. The volunteers were asked to perform these postures around 1.5 m in front of the mmWave device, but we did not limit the specific user position. While we are collecting mmWave data, we use a Microsoft Kinect (Zhang, 2012) to record the ground truth 3D coordinates of 17 skeletal joints (i.e., the same set of joints evaluated in WiPose (Jiang et al., 2020)) with a rate of 30 frames per second. The Kinect was placed behind the mmWave device as shown in Fig. 1(a) and extracts 17 skeletal joints with vision-based techniques (Microsoft, 2021). Around 640 samples (i.e., mmWave data frames with ground truth skeleton coordinates) were collected per subject per posture. We split the whole dataset (mmWave samples with ground truth joint coordinates) into a training set and a testing set, by using a ratio of 4:1. The deep learning model of *mPose* is trained with the samples and ground truth in the training set.

**Evaluation Metrics.** we define the joint localization errors in the 3D space as the projection of the distance between each joint's predicted coordinates and the ground truth on the three axes, depth, azimuth, and elevation. Furthermore, we use the 3D Euclidean distance in millimeters (mm) to quantify the overall joint localization error across the three axes.

**Table 1**  
Overall performance of joint location error.

Location Error	3D Joint	Depth	Azimuth	Elevation
WiPose (Jiang et al., 2020)	36.7 mm	–	–	–
mm-Pose (Sengupta et al., 2020)	–	32.0 mm	75.0 mm	27.0 mm
RF-Pose3D (Zhao, Tian, et al., 2018)	76.7 mm	42.0 mm	49.0 mm	40.0 mm
<b>mPose</b>	<b>30.1 mm</b>	<b>10.7 mm</b>	<b>16.8 mm</b>	<b>15.0 mm</b>



**Fig. 7.** Reconstructed skeleton postures and ground-truth of four different posture.

## 6.2. Performance of 3D posture reconstruction

**Overall Performance.** We first evaluate the overall performance of *mPose* using data collected in the three room environments for 3D skeleton posture reconstruction. Table 1 shows the average joint localization error across all postures of the 7 subjects. The results show that *mPose* has low errors for all metrics. Specifically, the 3D joint localization error of *mPose* is 30.1 mm, and the errors in depth, azimuth, elevation are 10.7 mm, 16.8 mm, 15.0 mm, respectively. Compared to *mm-Pose* (Sengupta et al., 2020), which uses two mmWave devices for posture reconstruction, our system can achieve 21.3 mm, 58.2 mm, and 12.0 mm improvement in depth, azimuth, elevation with only a single mmWave device. In addition, *mPose* shows 6.6 mm improvement in joint localization compared to *WiPose* (Jiang et al., 2020), without using spatially distributed antennas/devices. Furthermore, *mPose* has less than half of the 3D joint localization error compared to *RF-Pose3D* (Zhao, Tian, et al., 2018), which uses a customized FMCW radio for posture reconstruction. These results demonstrate that *mPose* outperforms the state-of-the-art skeleton reconstruction schemes.

Fig. 7 illustrates the differences between the skeleton postures reconstructed by *mPose* and the ground truth. The four postures are: (a) waving hands, (b) lifting the right hand to the front for 180 degrees, (c) lifting the left hand to the front for 180 degrees, (d) lifting the right leg for 90 degrees, respectively. We can find that the reconstructed 3D skeletons are almost the same as the ground truth joint coordinates estimated with Kinect. Fig. 8 shows the 3D joint localization error for individual joints. It shows that our system can effectively localize most of the joints with errors lower than 25 mm, except for joint 6 (elbow left), joint 7 (wrist left), joint 9 (elbow right), and joint 10 (wrist right), since these joints are very close to each other in the arm. The results demonstrate the effectiveness of *mPose*.

**Impact of Subjects.** We further evaluate the generalizability of *mPose* to different subjects. Specifically, we train and test the deep learning model with data collected from each of the 5 subjects. Fig. 9(a) shows the joint reconstruction error of individual subjects. We can find that the 3D Euclidean errors are all below 32 mm, which indicate the robustness of *mPose* on different subjects. Furthermore, we observe that *mPose* has low errors in depth, azimuth, elevation, which are below 14 mm, 20 mm, and 19 mm, respectively, indicating that *mPose* is general and can be applied to subjects that have various heights and body sizes.

**Impact of Location.** Due to the short wavelength, the mmWave signals attenuate rapidly as the propagation distance increases, which introduces uncertainty to our system. Hence, we evaluate the robustness of *mPose* under 3 different distances (i.e., 1.5 m, 2.5 m, and 3.5 m) between the target subject and mmWave device, with the deep learning model trained with the data collected at




Fig. 8. Performance on reconstructing individual joints.




Fig. 9. Performance under different impacting factors.

1.5 m. Fig. 9(b) shows the joint localization errors of *mPose* under the three distances. We can see that *mPose* achieves the minimum joint reconstruction error at the distance of 1.5 m, where the training data are collected. We have a similar observation for the errors in depth, azimuth, and elevation. For the other distances, the errors increase a little bit, with a joint localization error of 47.4 mm at 2.5 m. Even the distance increases to 3.5 m, the joint localization error is still below 60 mm. This is because our environment interference removal algorithm can dynamically detect the subject's location and remove the impact of distance via interpolation, so as to ensure acceptable performance for the skeleton posture reconstruction at different distances.

**Impact of Environments.** To examine the performance of *mPose* in different environments, we take turns to train and test the deep learning model with data collected in each of the three environments (i.e., lab, office, apartment). As shown in Fig. 9(c), we can see that the joint localization errors are below 33 mm for all environments. The errors in depth, azimuth, and elevation are below 17.3 mm, 18.2 mm, and 22.9 mm, respectively. In addition, we find that *mPose* has relatively higher errors in the apartment. Such an effect is due to the more complex room layout and furniture placement of the apartment, which could introduce more complex multi-path effects of mmWave signals. But even in this case, *mPose* can still achieve a low joint localization error of 32.7 mm. These results demonstrate that *mPose* can achieve satisfactory performance on 3D skeleton reconstruction in different environments.

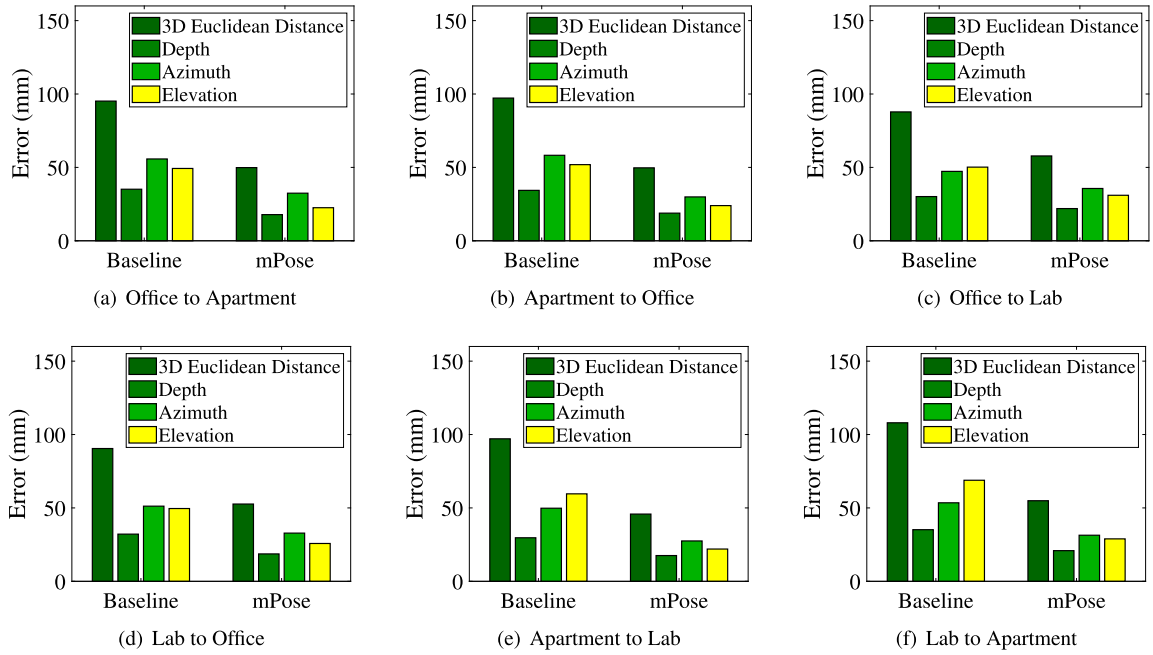


Fig. 10. Average joint localization errors for cross-environment 3D skeleton posture construction.

**Impact of Non-Line-of-Sight Conditions.** Different from vision-based approaches, mmWave signals can penetrate the obstacle between the mmWave device and the target subject. Hence, we evaluate the performance of *mPose* under the occlusions of metal and wooden objects, simulating object blocking in real-world scenarios (e.g., by home devices or furniture). In the experiment, we place an iron block and a wooden block in front of the mmWave device with a distance of 10 mm, which blocks the LOS between the subject and the mmWave device. Fig. 9(d) shows the 3D joint reconstruction error of *mPose* in the LOS and the NLOS scenes. We can see that although the joint localization errors under the NLOS scenarios are larger than that under the LOS scenarios, *mPose* can achieve 62.4 mm joint localization error when a wooden block placed in front of the mmWave device. In addition, we find that even when the LOS is blocked by iron block, where the mmWave signals are hard to penetrate, *mPose* can still achieve 63.9 mm. These results indicate that our system can work in NLOS scenarios with acceptable performance.

### 6.3. Performance of environment-independent posture reconstruction

We next evaluated the transferability of *mPose* across environments. Particularly, we trained the deep learning model with mmWave data from one environment and evaluated it with data collected in another environment, without additional training. Such an evaluation was conducted on every pair of the three environments (i.e., office, lab, and apartment). To show the effectiveness of the proposed environmental artifact removal (Section 5.3), we also trained and tested a deep learning model without applying the environmental artifact removal algorithm (i.e., the baseline model). Additionally, we show the overall joint localization error that is calculated by averaging the errors of all joints and all users. For the office and apartment environments, as shown in Fig. 10(a) and (b), we find that *mPose* achieves low joint localization errors of 49.8 mm and 49.7 mm under the two training–testing settings. Compared to the baseline model, we find that the proposed environmental artifact removal algorithm greatly reduces the posture reconstruction errors, with over 45.3 mm and 47.5 mm improvements for the two settings. The algorithm also reduces the reconstruction errors of depth, azimuth, and elevation, showing its effectiveness to improve the reconstruction in all dimensions. Similarly, as shown in Fig. 10(c) and (d), *mPose* shows lower joint localization errors with the environmental artifact removal algorithm for the office and lab environments, with over 33 mm and 31 mm improvements compared to the baseline model for the two training–testing settings. Similar observations can be found in Fig. 10(e) and (f), our system can achieve relatively lower 3D joint localization errors of 68.1 mm and 69.3 mm for the domain adaptation between the apartment and the lab. In general, the results demonstrate that the proposed environmental artifact removal algorithm is highly effective in improving the transferability of *mPose* across different environments.

### 6.4. Performance of subject-independent posture reconstruction

In real-world scenarios, the posture reconstruction model is usually trained with the dataset from limited users, which makes the model fail to reconstruct posture of users do not involve in the training. To evaluate *mPose* in the subject-independent scenarios, we trained the deep learning model with data collected from one subject and adapt the model to another subject using the domain




Fig. 11. Average joint localization errors for cross-subject 3D skeleton posture construction.

adaptation technique described in Section 5.4. For comparison, we also examined the baseline deep learning model without domain adaptation (i.e., the baseline model). We repeated such an evaluation for all subject pairs and averaged the errors. Fig. 11(a) shows the reconstruction errors of *mPose* when training and testing on different subjects in the lab environment. We can observe that *mPose* has 59.6 mm average joint localization error when using the proposed domain adaptation method, which is improved by 47.7 mm compared to the baseline model. We also find that the domain adaptation method helps to achieve better performance on individual dimensions in the 3D space, with 10.3 mm, 24.7 mm, and 34.1 mm improvements in the dimension of depth, azimuth, and elevation, respectively. Similar results can be found in the office environment as shown in Fig. 11(b), with 53.2 mm lower joint localization error after applying the domain adaptation method. The results demonstrate that the proposed domain adaptation method can help to achieve reliable posture reconstruction across subjects with unique body shapes.

## 7. Conclusion

In this paper, we proposed *mPose* which can continuously track a user's 3D skeleton postures using mmWave radars. The system can be deployed on a single portable COTS mmWave device while achieving a high joint localization accuracy. Through dynamically tracking the user's relative position to the system, *mPose* extracts spatial features associated with the user in a 3D contour, which removes the impacts of environmental factors. A deep regression model based on convolutional neural network is designed to map the extracted spatial features into the skeletal joint coordinates in a 3D space. With the designed domain adaptation method, *mPose* removes subject-specific characteristics entangled in the spatial features and enables robust posture reconstruction across users. Experimental results demonstrated that *mPose* can accurately reconstruct full-body 3D skeleton posture, achieving at least 18% lower prediction errors than existing device-free skeleton posture reconstruction methods with a single mmWave device. Additionally, the experiments confirm the robustness and generalizability of *mPose* with different environments and different users. We believe that our portable single-mmWave-device solution could enable various emerging applications, including immersive augmented reality or virtual reality, mobile healthcare, and pervasive security monitoring.

## CRediT authorship contribution statement

**Cong Shi:** Methodology, Software, Data Curation, Writing – original draft. **Li Lu:** Methodology, Writing – review & editing. **Jian Liu:** Methodology, Writing – review & editing. **Yan Wang:** Methodology, Writing – review & editing. **Yingying Chen:** Conceptualization, Supervision, Project administration, Writing – review & editing. **Jiadi Yu:** Writing – review & editing.

## Declaration of competing interest

We declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was partially supported by the National Science Foundation Grants CCF-2000480, CCF-2028858, CCF-2028876, CNS-2120276, CNS-2120396, ECCS-2033433, ECCS-2132106, and the Grant of National Natural Science Foundation of China No. 62172277. We thank Song Yang and Zhenzhe Lin in their efforts on conducting experiments and analyzing data.



## References

- Adib, F., Hsu, C.-Y., Mao, H., Katabi, D., & Durand, F. (2015). Capturing the human figure through a wall. *ACM Transactions on Graphics*, 34(6), 1–13.
- Boulay, B., Bremond, F., & Thonnat, M. (2003). Human posture recognition in video sequence. In *International workshop on visual surveillance and performance evaluation of tracking and surveillance*.
- Cao, Z., Hidalgo, G., Simon, T., Wei, S.-E., & Sheikh, Y. (2018). Openpose: realtime multi-person 2D pose estimation using part affinity fields. *arXiv Preprint arXiv:1812.08008*.
- Caserman, P., Garcia-Agundez, A., Konrad, R., Göbel, S., & Steinmetz, R. (2019). Real-time body tracking in virtual reality using a Vive tracker. *Virtual Reality*, 23(2), 155–168.
- Chen, Y., & Xue, Y. (2015). A deep learning approach to human activity recognition based on single accelerometer. In *Proceedings of International Conference on Systems, Man, and Cybernetics* (p. 1488–1492).
- Feng, L., Li, Z., Liu, C., Chen, X., Yin, X., & Fang, D. (2020). SitR: Sitting posture recognition using RF signals. *IEEE Internet of Things Journal*, 7(12), 11492–11504.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., et al. (2016). Domain-adversarial training of neural networks. *The Journal of Machine Learning Research (JMLR)*, 17(1), 2096–2030.
- Guo, H., Yu, Y., Ding, Q., & Skitmore, M. (2018). Image-and-skeleton-based parameterized approach to real-time identification of construction workers' unsafe behaviors. *Journal of Construction Engineering and Management*, 144(6), Article 04018042.
- Ho, E. S., Chan, J. C., Chan, D. C., Shum, H. P., Cheung, Y.-m., & Yuen, P. C. (2016). Improving posture classification accuracy for depth sensor-based human activity monitoring in smart environments. *Computer Vision and Image Understanding*, 148, 97–110.
- (2021). VIVE - Tracker. HTV. <https://www.vive.com/us/vive-tracker/>.
- Iovescu, C., & Rao, S. (2017). The fundamentals of millimeter wave sensors. Texas Instruments.
- Iwasawa, S., Ohya, J., Takahashi, K., Sakaguchi, T., Ebihara, K., & Morishima, S. (2000). Human body postures from trinocular camera images. In *Proceedings fourth ieee international conference on automatic face and gesture recognition (Cat. No. PR00580)* (pp. 326–331). <http://dx.doi.org/10.1109/AFGR.2000.840654>.
- Jiang, W., Xue, H., Miao, C., Wang, S., Lin, S., Tian, C., et al. (2020). Towards 3D human pose construction using wifi. In *Proceedings of the 26th annual international conference on mobile computing and networking* (p. 1–14).
- Jiang, W., & Yin, Z. (2015). Human activity recognition using wearable sensors by deep convolutional neural networks. In *Proceedings of the international conference on multimedia* (p. 1307–1310).
- Lien, J., Gillian, N., Karagozler, M. E., Amihoud, P., Schwesig, C., Olson, E., et al. (2016). Soli: Ubiquitous gesture sensing with millimeter wave radar. *ACM Transactions on Graphics*, 35(4), 1–19.
- Liu, J., Wang, Y., Chen, Y., Yang, J., Chen, X., & Cheng, J. (2015). Tracking vital signs during sleep leveraging off-the-shelf wifi. In *Proceedings of the 16th ACM international symposium on mobile ad hoc networking and computing* (p. 267–276).
- Mehrizi, R., Peng, X., Xu, X., Zhang, S., Metaxas, D., & Li, K. (2018). A computer vision based method for 3D posture estimation of symmetrical lifting. *Journal of Biomechanics*, (69), 40–46.
- (2021). Kinect. Microsoft. <https://en.wikipedia.org/wiki/Kinect>.
- Mueller, F., Davis, M., Bernard, F., Sotnychenko, O., Verschoor, M., Otaduy, M. A., et al. (2019). Real-time pose and shape reconstruction of two interacting hands with a single depth camera. *ACM Transactions on Graphics*, 38(4), 1–13.
- Papandreou, G., Zhu, T., Kanazawa, N., Toshev, A., Tompson, J., Bregler, C., et al. (2017). Towards accurate multi-person pose estimation in the wild. *arXiv:1701.01779*.
- (Don't) hold the phone: new features coming to pixel 4. (2019). <https://www.blog.google/products/pixel/new-features-pixel4/>.
- Postolache, O., Girão, P. S., Madeira, R. N., & Postolache, G. (2010). Microwave FMCW Doppler radar implementation for in-house pervasive health care system. In *2010 IEEE international workshop on medical measurements and applications* (pp. 47–52). IEEE.
- Quwaider, M., & Biswas, S. (2008). Body posture identification using hidden Markov model with a wearable sensor network. In *Proceedings of the international conference on body area networks* (p. 1–8).
- Rao, S. (2017). Introduction to mmWave sensing: FMCW radars. Texas Instruments (TI) MmWave Training Series.
- Ren, Z., Yuan, J., Meng, J., & Zhang, Z. (2013). Robust part-based hand gesture recognition using kinect sensor. *IEEE Transactions on Multimedia*, 15(5), 1110–1120.
- Schwarz, L. A., Mkhitarany, A., Mateus, D., & Navab, N. (2012). Human skeleton tracking from depth data using geodesic distances and optical flow. *Elsevier Image and Vision Computing*, 30(3), 217–226.
- Sengupta, A., Jin, F., Zhang, R., & Cao, S. (2020). Mm-pose: Real-time human skeletal posture estimation using mmwave radars and CNNs. *IEEE Sensors Journal*.
- Shen, S., Wang, H., & Roy Choudhury, R. (2016). I am a smartwatch and i can track my user's arm. In *Proceedings of the 14th annual international conference on mobile systems, applications, and services* (p. 85–96).
- Stove, A. G. (1992). Linear FMCW radar techniques. 139, In *IEEE proceedings F (radar and signal processing)* (5), (pp. 343–350). IET.
- Tompson, J., Jain, A., LeCun, Y., & Bregler, C. (2014). Joint training of a convolutional network and a graphical model for human pose estimation. In *Proceedings of the 27th international conference on neural information processing systems* (p. 1799–1807).
- Valero, E., Sivanathan, A., Bosché, F., & Abdel-Wahab, M. (2017). Analysis of construction trade worker body motions using a wearable and wireless motion sensor network. *Automation in Construction*, 83, 48–55.
- VICON (2019). Motion capture systems | VICON. <https://www.vicon.com/>.
- Wang, Y., Liu, J., Chen, Y., Gruteser, M., Yang, J., & Liu, H. (2014). E-eyes: device-free location-oriented activity identification using fine-grained wifi signatures. In *Proceedings of the 20th annual international conference on mobile computing and networking* (p. 617–628).
- Wang, F., Zhou, S., Panev, S., Han, J., & Huang, D. (2019). Person-in-WiFi: Fine-grained person perception using WiFi. In *Proceedings of the IEEE/CVF international conference on computer vision* (p. 5452–5461).
- Wikipedia (2019). Huber loss. [https://en.wikipedia.org/wiki/Huber\\_loss](https://en.wikipedia.org/wiki/Huber_loss).
- Xu, X., Gao, H., Yu, J., Chen, Y., Zhu, Y., Xue, G., et al. (2017). ER: Early recognition of inattentive driving leveraging audio devices on smartphones. In *IEEE conference on computer communications* (p. 1–9).
- Yang, X., Liu, J., Chen, Y., Guo, X., & Xie, Y. (2020). MU-ID: Multi-user Identification Through Gaits Using Millimeter Wave Radios. In *Proceedings of international conference on computer communications* (p. 2589–2598).
- Yang, Z., Pathak, P. H., Zeng, Y., Liran, X., & Mohapatra, P. (2017). Vital sign and sleep monitoring using millimeter wave. *ACM Transactions on Sensor Networks*, 13(2), 1–32.
- Zhang, Z. (2012). Microsoft kinect sensor and its effect. *IEEE Multimedia*, 19(2), 4–10.
- Zhao, M., Li, T., Abu Alsheikh, M., Tian, Y., Zhao, H., Torralba, A., et al. (2018). Through-wall human pose estimation using radio signals. In *Proceedings of the conference on computer vision and pattern recognition* (p. 7356–7365).
- Zhao, M., Tian, Y., Zhao, H., Alsheikh, M. A., Li, T., Hristov, R., et al. (2018). RF-based 3D skeletons. In *Proceedings of the special interest group on data communication* (p. 267–281).