

Gesture Recognition on Deformable Objects Using Millimeter-Wave Radar

Nuwan T Attygalle

Louvain Research Institute in
Management and Organizations
Université Catholique de Louvain
Louvain-la-neuve, Belgium
nuwan.attygalle@uclouvain.be

Matjaž Kljun

University of Primorska
Koper, Slovenia
Stellenbosch University
Stellenbosch, South Africa
matjaz.kljun@upr.si

Klen Čopič Pucihar

University of Primorska
Koper, Slovenia
Stellenbosch University
Stellenbosch, South Africa
klen.copic@famnit.upr.si




Figure 1: Wrist-worn millimetre-wave (mmWave) radar sensor (left); experimental setup with a participant seated, holding a plush toy while performing instructed gestures displayed on a laptop screen (center); example set of gestures performed on the plush toy (right).

Abstract

Although deformable objects are not typically designed for digital interaction, they offer a largely unexplored potential—any such object could be repurposed as a medium for controlling digital content. While existing approaches embed sensors into deformable objects to enable interaction, this limits scalability and practicality of such systems. An alternative is to perform gesture recognition on deformable objects using a wrist-worn radar sensor. However, when analysing reflected radar signals it is difficult to separate reflections originating from the continuous deformations of the object shape and those from the user’s hand and fingers. Additionally, the continuous shape changes of deformable objects introduce changes in radar cross-section, affecting signal variability. Furthermore, user ergonomics—such as variations in hand size, finger dexterity, and strength—are likely to influence the degree of object deformation during interaction. In this paper, we explore whether radar sensing can be used for robust gesture detection on deformable objects, focusing on how well does a system generalize to previously unseen users and what can we do to improve such generalisability. In pursuit of this goal, we record a dataset of 4.3k labelled gestures with Google Soli millimeter-wave radar sensor on a plush

toy and demonstrates robust classification performance, achieving accuracy of up to 90% on a five-gesture set. Furthermore, we investigate model generalizability and show that transfer learning improves recognition for previously unseen users, yielding performance gains of up to 20%. These findings highlight the potential of radar-based sensing for spontaneous and practical interaction with deformable objects.

CCS Concepts

• **Human-centered computing** → **Gestural input**; **Graphical user interfaces**; *Interactive systems and tools*; • **Computing methodologies** → *Model development and analysis*; • **Software and its engineering** → Runtime environments; • **Hardware** → Radio frequency and wireless interconnect.

Keywords

Gestures on Deformable, mmWave radar, Deformable objects

ACM Reference Format:

Nuwan T Attygalle, Matjaž Kljun, and Klen Čopič Pucihar. 2025. Gesture Recognition on Deformable Objects Using Millimeter-Wave Radar. In *Companion Proceedings of the 17th ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS Companion '25)*, June 23–27, 2025, Trier, Germany. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3731406.3734979>

1 Introduction

Most deformable objects made from fabric, textiles, foam or silicon, such as for example plush toys, clothing, upholstery, seat cushions, elastic bands, and rubber or silicone mats are not designed with the intention to support digital interaction. Nevertheless, such

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

EICS Companion '25, Trier, Germany

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1866-3/25/06

<https://doi.org/10.1145/3731406.3734979>

deformable objects present a largely untapped opportunity for interaction with digital content. For instance, consider a textile-based gaming input, where squeezing or twisting a soft toy made from fabric and foam could be used to execute in-game actions. Another example could consider interacting with soft therapy devices, such as for example pressing, stretching, or squeezing deformable rehabilitation tools to provide input for physical therapy tracking and gamification of rehabilitation exercises.

A common approach to enabling such interaction is to embed computational capabilities into objects, making them “smart” and allowing users to control aspects of their digital environment. An alternative would be to detect gestures performed directly on deformable objects without instrumenting them. As any object can become an interface, we significantly expand interaction possibilities.

However, gesture recognition on deformable objects presents distinct challenges, particularly when these objects are not instrumented with sensors. Traditional vision-based gesture detection systems require a direct line of sight and have a hard time detecting gestures when occlusion happens [1, 16, 46]. This limitation significantly constrains interaction possibilities, especially when users grip, bend, or manipulate the object, which in turn results in partial occlusion of their hands and/or fingers.

Unlike mid-air gestures in general [12] or with radars [33, 36, 44] or on-object [1] gesture recognition, deformation-based gestures [11, 15] rely on capturing shape changes of the object, material flexibility, and force dynamics, allowing users to interact with digital systems through bending, stretching, squeezing, or twisting deformable objects [11, 14]. Common approaches for detecting object deformation involve vision-based approaches that use RGB and RGB-D cameras. Again, vision-based techniques face challenges such as occlusions, lighting conditions, and high computational costs. To overcome these issues, sensor-based techniques—including strain gauges, piezoelectric sensors, capacitive sensing, and soft resistive materials—have been developed to directly measure deformations on flexible surfaces and objects. However, instrumenting objects with additional sensors is not practical thus limiting the applicability of such solutions.

Radar-based sensing presents a promising alternative for spontaneous interaction with deformable objects and surfaces, particularly in scenarios where gestures are partially occluded. Unlike optical methods, radar technology operates without requiring a direct line of sight, as electromagnetic waves can penetrate non-conductive materials. This paper explores if this inherent capability of radar enables a robust gesture detection on deformable objects, despite the unique challenges posed by object’s continuous shape deformations and object’s material properties.

2 Problem statement and contributions

When analysing reflected radar signals of gestures performed on deformable objects, it is difficult to separate reflections originating from the continuous deformations of the object shape and those from the user’s hand and fingers. Theoretically one could separate the object, and hand and finger based on range, however as the object and hand are very close to each other, this is difficult to do with

analytical methods, especially when considering dynamic environments which are common for gesture interaction on deformable objects. Therefore, instead of segmenting the hand and the object we opt for combined detection that includes hand movements and object deformations.

Additionally, the continuous shape changes of deformable objects introduce changes in radar cross-section, affecting signal variability. For instance, flat surfaces tend to produce strong reflections when aligned perpendicularly to the radar wave, whereas tilted surfaces redirect reflections away from the receiver, leading to signal loss. However, as in the case of deformation-based gestures, the deformation of the object itself may contain valuable information about the gesture being performed. Yet, user ergonomics—such as variations in hand size, finger dexterity, and applied force—are likely to significantly influence the degree of object deformation during interaction. This variability results in diverse deformation patterns, further compounding inconsistencies in radar signal reflections. As a consequence, gesture recordings exhibit high intra- and inter-user variability, making it challenging to develop robust gesture recognition models that generalize across unseen users.

This opens up several research questions we address in this paper: **RQ1**: Can radar sensing be used for robust gesture detection on deformable objects? If so, **RQ2**: How well do recognition models generalize to previously unseen users? And **RQ3**: What strategies can be employed to improve model adaptability and robustness across unseen users?

To answer these questions, we designed, implemented, and evaluated a gesture recognition system based on Google Soli’s mmWave radar sensing technology. We show that our system can detect different gestures on and behind deformable objects, revealing its potential for spontaneous interactions with everyday deformable objects. The contributions of this paper are:

- A publicly available data set of 4300 labelled gestures executed on a plush toy using a wrist-mounted mmWave Google Soli radar sensor (<https://gitlab.com/hicuplab/godo-soli>).
- Gesture recognition system that can enable robust gesture detection on deformable objects (performance of up to 95% on a five-gesture set).
- A comprehensive set of experiments involving the training and evaluation of 20 models, showing how transfer learning could be used to improve model generalisability to new users, with performance gains of up to 20%.

3 Related work

Everyday objects are typically not designed with built-in interactive capabilities. However, this does not prevent users from interacting with them in meaningful ways. For example, cuddly teddy bears are widely used for emotional comfort across all stages of life [37]. Their visual appeal, olfactory characteristics, softness, and kinaesthetic properties contribute to their comforting effect [37]. These factors make teddy bears more than just toys—they are emotionally significant objects for users [31, 37].

Given this strong emotional attachment, embedding sensors into an existing plush toy is often impractical as it may alter its softness and kinaesthetic feel, which are unique to plush toys compared to rigid objects. Enabling interaction without physically modifying the

toy would preserve its natural comforting properties. Communicating with a user's favourite teddy bear through gesture recognition creates a wholesome and intuitive experience.

Since mmWave radar signals can penetrate non-conductive materials [46] and recognize subtle hand gestures, we explore the use of a wrist-mounted mmWave radar sensor for detecting micro-gestures on a plush toy. This approach allows for seamless interaction without embedding any sensors into the object itself.

3.1 Sensing Technologies for Deformable Object Interaction

Sensing interactions on deformable objects has been explored using various sensor technologies, including capacitive sensors[20], conductive wire-based sensors[23], impedance-based sensors[43], and flex sensors[5], all requiring embedded sensors for functionality.

Maiolino et al. [20] introduced capacitive sensors for non-flat surfaces, enhancing humanoid robot interactions[27, 28]. SmartSleeve[23] used conductive wires to enable squeezing, swiping, and touching on garments. FoamIn[43] applied impedance sensing to capture gestures on dynamically deforming foam.

Vision-based methods have also been explored. Mahdikhani and Ebrahimzadeh [19] used colored gloves and cameras to track 3D object deformations, while GraspUI[29] tracked grasp gestures on deformable objects using dual-camera setups.

Alternative sensing methods include RaTIn[7], which identifies objects via radar and reflective markers, and Counterpoint[6], which uses radar and Leap Motion for mixed-scale AR gestures but does not address deformable materials. Touché[25] captures environmental gestures with capacitive sensing but requires conductive surfaces, while FuwaFuwa[10] embeds photoreflexive sensors into soft toys. ViBand[13] detects bio-acoustic signals through smart-watch accelerometers rather than object deformation. Wearable sensing has also been investigated, such as ear gesture recognition using IMUs in earphones[26].

Existing methods either embed sensors into objects or rely on vision-based systems with occlusion and privacy limitations. Therefore, we propose using radar sensors for micro-gesture detection on deformable objects, leveraging radar's ability to penetrate non-conductive materials such as plush toys. To the best of our knowledge, gesture sensing on deformable objects with radars has not yet been explored.

3.2 Micro-gesture Sensing Technologies

Micro-gestures are detected using vision-based tracking, wearable sensors, or radar systems. Vision-based systems like Kinect [4] and Leap Motion [35] are accurate but sensitive to lighting and occlusion. Wearable sEMG sensors [38] and Tap Strap 2 [21] enable continuous recognition but require user contact and may be cumbersome. On the other hand radar-based sensing [46] provides a contactless, privacy-preserving alternative, unaffected by lighting or field of view [34].

Machine learning advancements, including CNNs [3], RNNs [45, 46], WE-KNN [17], and model-based approaches [36], enhance micro-gesture detection accuracy, enabling real-time interpretation and robust differentiation of fine hand movements. These advances

also contributed to the development of radar-based gesture recognition.

Previous work has investigated radar-based gesture recognition in various contexts, but none have specifically addressed deformable objects. RadarHand[8] explored wrist-worn Soli radar to capture on-skin touch-based gestures, while The Wearable Radar[16] examined gesture recognition through fabric materials. Solids-on-Soli[46] extended this concept by analyzing radar reflections across 73 everyday materials. The closest related work, No Interface, No Problem[1], studied gesture recognition on a picture frame using a wrist-mounted Soli sensor. However, none of these studies specifically investigated gestures performed on deformable objects, leaving a significant gap that our work aims to address.

3.3 Wrist-mounted sensors

Prior research has demonstrated the effectiveness of wrist-mounted sensors, particularly cameras, for enabling gesture and touch-based interactions. For example, Ohnishi et al. [22] used a wrist-mounted camera to recognize daily activities involving large hand movements and object interactions. Similarly, FingerTrak [9] employed thermal cameras mounted on a wristband to continuously capture 3D hand pose for applications such as sign language recognition. WristCam [2] introduced a wearable camera for tracking hand trajectory gestures to support human-robot interaction, while GestureWrist [24] enabled unobtrusive gesture input using wrist-worn sensing. These studies collectively motivate the wrist as a practical and effective location for sensing fine-grained hand gestures due to its proximity to the hand and ability to maintain a stable sensing perspective. Inspired by this body of work, we mount a Soli radar sensor on the wrist to classify on-body gestures performed on a handheld plush toy. Radar sensing at the wrist offers a compact, low-power alternative to vision-based systems, while preserving the benefits of continuous, unobtrusive interaction capture.

4 Gesture recognition experiments

This section outlines the method used in the gesture recognition experiments, beginning with the gesture set, data collection, and preprocessing steps. It then details the data partitioning strategy, model architecture, training procedure, and domain adaptation approach.

4.1 Gesture set

We selected five gestures (Figure 1 right) for interaction with deformable objects: Thumb [1], Thumb Joint, Scratch, Pinch [41], and Squeeze [40] (see Figure 1 right). The selection was informed by a review of micro-gestures previously detected using radar systems[18, 42], with a particular focus on on-object micro-gestures [1].

Additionally, we conducted a heuristic analysis considering both the sensing capabilities of radar systems and gesture ergonomics. The authors experimented with various gestures, prioritizing those that could be performed comfortably and consistently on deformable objects. Only gestures that could be executed with ease and were also considered as radar friendly (e.g. exhibiting substantial amount of motion in range) were included in the final set.

The selected gestures fall into two distinct categories based on the extent of object deformation: (i) Gestures causing minimal

object deformation (Thumb, Thumb Joint, and Scratch); (ii) Gestures causing large object deformation (Pinch and Squeeze). The first group of gestures aligns with findings from grasping micro-gesture elicitation studies. However, the second group—gestures involving large deformations—was not observed in previous studies. This is likely because prior elicitation studies were conducted on rigid objects [30], which inherently restricted the feasibility of gestures that involve substantial object deformations.

4.2 Data Collection and Pre-Processing

With the help of ten participants (6 male, 4 female, aged 23–50) we recorded 4.3k labelled instances (5 gestures, 10 users, 20 gesture repetitions, 2–5 sessions) following the same procedure as was followed in [39]. In each session, 20 repetitions were made of each gesture and sessions were spaced out by a minimum of 1 hour. Each gesture recording is 1-second long and was executed on a plush toy which was held by participants in the lap whilst sitting on the chair. All instructions were provided on a laptop screen that was placed on a table in front of participants (see Figure 1 centre).

The screen showed an animated image of the gesture to be executed. After a beep sound was played, participants had to execute the gesture within a 1-second time window. This was repeated 20 times for each gesture. The order of gestures was randomised and after each round, the sensor was reset (the clutter map of the radar sensor was rebuilt). Participants were also asked to repeat the gesture if the one they executed did not match the gesture shown on the image (e.g., the user executed the scratch gesture instead of the thumb gesture).

The Soli sensor was configured to record on 4 channels at 1000Hz with the adaptive clutter filter disabled. The recorded files were pre-processed at a dB range setting of $[-32, 0]$, downsampled to 100Hz, and greyscaled, generating 4 Range-Doppler images per frame. Each Range-Doppler image corresponds to $T_X - R_X$ antenna pairs and consists of 32 range bins and 32 velocity bins (i.e. 32×32 pixels). The intensity of the image coincides with the radar-cross section of the target. Further, negative velocity corresponds to the motion towards the soli sensor, while the positive velocity corresponds to the motion away from the sensor.

Each gesture recording is 1 second long, thus we have 100 frames per recording. To reduce the number of images, we flattened each Range-Doppler image to 1024 bins and stack them on top of each other, resulting in a spectrogram image of 1024×100 pixels (Figure 2) for each channel. Finally, we stack each antenna configuration on top of each other creating a spectrogram image of 400×1024 . Making spectrograms in this way allowed us to preserve phase variations along with range, velocity and radar-cross section details of the target.




Figure 2: A spectrogram image for a Thumb gesture, channel 0, pre-processed with the dB range $[-32, 0]$. The image is generated by flattening a sequence of 100 Range-Doppler images. Each row represents one Range-Doppler image of $32 \times 32 = 1024$ px.

4.3 Data partitioning

All experiments are done on two different datasets, namely FULL-DATASET and REDUCED-DATASET. Both datasets are balanced on users and gestures. This is crucial for Leave-one-user-out evaluations. However, due to the fact that not all users participated in the same number of sessions, we could only use 2000 recordings ($40 \text{ repetitions} * 5 \text{ gestures} * 10 \text{ users}$) in the full-dataset. In the reduced dataset, we only use 750 recordings ($15 \text{ repetitions} * 5 \text{ gestures} * 10 \text{ users}$). This decision was made due to the high cost of data capturing as well as higher possibility for gain when attempting domain adaptation.

We run two experiments, All-users evaluation and Leave-one-user out evaluation. In All-users evaluation 10% of data was used for test partition (200 recordings (40 per gesture) and 75 recordings (15 per gesture) for full-dataset and reduced-dataset respectively). The remainder (1800 and 675 for full-dataset and reduced-dataset respectively) was further segmented into train(70%) and validation partitions(30%). Leave-one-user out was done in a standard way, we train the model once per user, leaving one user out and using this user as our test partition. Other users form train(70%) and validation (30%) partitions.

4.4 Model architecture

We adopted the classifier architecture introduced in the *No-Interface-No-Problem* [1] due to its proven performance, lightweight design, and ease of implementation. The original model is a compact CNN-based network optimized for spectrogram inputs, with no recurrent (e.g., LSTM) layers, making it suitable for deployment on wearable devices. It has approximately 770K trainable parameters and a model size of just 8.9MB, which is considerably small compared to typical deep learning models [1]. In our implementation, we modified the input layer to accommodate the dimensions of our spectrogram data and adjusted the output layer to match our gesture class set.

The model is based on 2D CNN layers, and it has four Conv2D layers each containing 32, 64, 128, and 256 filters with each consisting (3×3) kernels with *LeakyReLU* activations on each Conv2D layer. Further, each Conv2D layer is followed by a MaxPool2D layer for downsizing the output of each Conv2D layer. After the Conv2D layers, it goes to flatten layer and finally put out to Dense layer with *Softmax* activation for 5 classes.

4.5 Model training and domain adaptation

We experiment with two different training scenarios, in the baseline scenario (BASELINE MODEL) the model weights were trained from scratch, whilst in the domain adaptation scenario the models were trained with transfer learning (TL MODEL). In both cases the models were trained for a maximum of 100 epochs. To avoid overfitting, we use an early stopping criteria (*patience* = 20). We also used learning rate plateau middleware to reduce the learning rate on plateau with *patience* = 3 and factor 0.5. We are always starting the learning rate on $1.e - 4$ and set the minimum learning rate to $1.e - 6$. We used the usual categorical cross entropy for the loss function since this is simple multi-class classification, and Adam optimizer for updating the model parameters calculated in backpropagation.

In the domain adaptation scenario with transfer learning, we employed a Solids-on-Soli model previously trained on a robot dataset by Čopić Pucihar et al. [46]. This model can classify 6 gestures, namely: swing left, swing right, away, towards, wiggle, and no action as depicted in Čopić Pucihar et al. [46]. The same sensor was used, with very similar configuration. We brought this Solids-on-Soli trained model into our experiment and as in a typical transfer learning scenario, we froze all the layers but removed the last layer and replaced it with a *Dense* layer with 5 outputs with a *Softmax* activation. Afterwards, we trained the model, but allowing to adjust only the neurons in the last layer. How different are the two gesture sets can be compared by observing Figure 1 right and gesture in Čopić Pucihar et al. [46]. How does this manifest in spectrogram representations of a radar signal can be observed on Figure 3 and Figure 4.

5 Results

5.1 All-users evaluation

Results in Table 1 show good performance of the BASELINE MODEL in full-dataset scenario ($acc = 89\%$, $AUC = 0.89$, $Precision = 0.9$, $Recall = 0.89$ and $F1 = 0.9$), yet this performance drastically falls in a reduced-dataset scenario ($acc = 29\%$, $AUC = 0.71$, $Precision = 0.39$, $Recall = 0.29$ and $F1 = 0.25$). In case of TL model, we only observe a small performance drop as the dataset size decreases. The accuracy drops from $acc = 90\%$, $AUC = 0.99$, $Precision = 0.9$, $Recall = 0.8$ and $F1 = 0.9$ in the full-dataset scenario to $acc = 81\%$, $AUC = 0.96$, $Precision = 0.82$, $Recall = 0.81$ and $F1 = 0.81$ in the reduced-dataset scenario. The same trends can also be observed in confusion matrices in Figure 5 left with the lowest accuracies reported in BASELINE MODEL for a reduced-dataset scenario in case of gestures G1-Thumb, G4-Pinch and G5-Squeeze. Interestingly, for gestures coming from the group causing large object deformations (gestures G4-Pinch and G5-Squeeze) transfer learning (e.g. TL MODEL) was not as effective (see top right confusion matrix in Figure 5 left). The confusion matrices also show TL MODEL has higher accuracy in full-dataset and reduced-dataset for all but one gesture. The confusion matrices Figure 5 left also show TL MODEL has higher accuracy in full-dataset and reduced-dataset for all but one gesture.

Model	Dataset	Performance	ACC	AUC	Precision	Recall	F1
Baseline	Full-dataset	0.94	0.89	0.99	0.9	0.89	0.90
TL	Full-dataset	0.945	0.90	0.99	0.90	0.90	0.90
Baseline	Reduced-dataset	0.5	0.29	0.71	0.39	0.29	0.25
TL	Reduced-dataset	0.885	0.81	0.96	0.82	0.81	0.81

Table 1: Performance of BASELINE MODEL and TC model when trained on all-users for full-dataset and reduced-dataset.

5.2 Leave-one-user-out (RQ2 and RQ3)

Similar trend can be observed as in All-users evaluation (subsection 5.1). Results in Table 2 show good user specific generalisation capabilities for the BASELINE MODEL in full-dataset scenario (median scores of: $acc = 73\%$, $AUC = 0.93$, $Precision = 0.76$, $Recall = 0.73$ and $F1 = 0.73$), yet this ability drastically falls in a

reduced-dataset scenario (median scores of: $acc = 38\%$, $AUC = 0.80$, $Precision = 0.52$, $Recall = 0.38$ and $F1 = 0.35$). In case of TL model (Table 3), we only observe a small performance drop in user specific generalisation capabilities as the dataset size decreases. The median scores of: $acc = 84\%$, $AUC = 0.98$, $Precision = 0.83$, $Recall = 0.84$ and $F1 = 0.82$ in the full-dataset scenario fall to median scores of: $acc = 73\%$, $AUC = 0.93$, $Precision = 0.76$, $Recall = 0.73$ and $F1 = 0.73$ in the reduced-dataset scenario.

To compare BASIC MODEL performance to TL MODEL performance, where performance is defined as $performance = (ACC + AUC)/2$ we run Paired-sample T-test Figure 6. Shapiro-Wilk test showed the data is normally distributed (see top table in Figure 6). The results in Figure 6 show significant improvement in performance for TL MODEL in a reduced-dataset scenario ($t(9) = 3.51$, $p = .007$). The mean performance increased from 63% (SD = 17.57%) to 82.2% (SD = 8.92%), 95% CI [6.79, 31.4]. The effect size was moderate (Cohen's $d = 1.11$), indicating a large improvement in performance. Despite also observing an increase of performance for full-dataset scenario, no significance was reached ($t(9) = 2.14$, $p = .06$).

The confusion matrices in Figure 5 right show the lowest accuracies in BASELINE MODEL for a reduced - dataset scenario in case of gestures G1-Thumb, G3-Scratch and G5-Squeeze. Unlike full-dataset experiment, where transfer learning only outperformed the BASELINE MODEL for gestures G1-Thumb, G3-Scratch and G5-Squeeze, in the reduced-dataset scenario it outperformed the baseline across all gesture classes (see top right confusion matrix in Figure 5 right). The confusion matrices also show TL MODEL has higher accuracy in full-dataset and reduced-dataset.

6 Discussion, limitations, and future work

We structure discussion based on the three research questions: Are radar-based gesture recognition models capable of robust gesture detection on deformable objects (RQ1); How well do these models generalize to previously unseen users (RQ2)? and whether domain adaptation is an effective strategy to improve generalization to unseen users (RQ3).

6.1 Robust gesture detection on deformable objects (RQ1)

Our results show deep neural networks can be used for robust gesture detection on deformable objects. This was the case despite being faced with the challenge of continuous shape changes of deformable objects, which introduce changes in radar cross-section. However, this was not the case as data became more scarce in reduced-dataset scenario. High volume data requirement is integral part of training deep neural network and besides computation cost and lack of explainability one of its key drawbacks.

An alternative approach would be to try and model interferences caused by object deformations through analytical methods, as was done in *RadarSense* [32, 33]. In this work the authors successfully removed interferences of the obstructed materials through inverse modelling in which the occluding material barrier became part of the radar antenna model producing clean spectrograms representations of gestures. These could be used to train less resource hungry machine learning algorithms, such as for example template matching. However, due to dynamic characteristic that are common in




Figure 3: Examples of generated spectrogram images (as explained in subsection 4.2) from the Range-Doppler images acquired from the Solids-on-soli [46] dataset. Each spectrogram contains 80 radar frames of a gesture execution, which are represented in 320 Range-Doppler images (80 per channel). These Range-Doppler images are flattened and stacked on top of each other.




Figure 4: Examples of spectrogram images from hand gestures on a plush toy. Each spectrogram contains 100 radar frames of a gesture execution, which are represented in 400 Range-Doppler images (100 per channel). These Range-Doppler images are flattened and stacked on top of each other.

gesture detection on deformable objects renders such an approach difficult and/or potentially impractical.

We acknowledge that incorporating explicit deformation information—via auxiliary sensors or deformation modeling—could further enhance recognition performance and interpretability. However, our aim in this work was to explore the feasibility of sensor-free gesture recognition on deformable objects using radar alone. In this paper, we address the small data size problem with domain adaptation of our model, discussed in subsections hereafter.

6.2 Model generalization to unseen users (RQ2)

Our experimentation with Leave-one-user-out show that as long as we have sufficient amount of data (Full-dataset scenario), our models generalise well to unseen users. This was the case despite gesture recordings exhibiting high intra- and inter-user variability, a result of variations in user ergonomics—such as variations in hand size, finger dexterity, and applied force— all of which may influence the degree of object deformation during interaction. This yet again shows that robust radar-based gesture recognition on deformable objects is possible. Our experimentation also showed it is essential

to have sufficient amount of data available. We again try to address the small data size problem with domain adaptation of our model, discussed in subsections hereafter.

6.3 Domain adaptation as a strategy to improve generalization (RQ3)

Domain adaptation through transfer learning proved an effective strategy for both all-users evaluation and leave-one-user-out evaluation. The effectiveness of domain adaptation was increased with the lack of data. This result is important because it shows how different the model used for transfer learning can actually be in this radar-based sensing context. For example, the model used in transfer learning was not only trained on a very different set of gestures, but these gestures were also performed by a robot system swinging the ball instead of a human subject performing gestures with a hand reducing the cost of such data collection.

6.4 Limitations and future work

We only experiment with 5 gestures executed on a plush toy. Also our gesture set is limited to bidirectional gestures. Therefore, future




Figure 5: Confusion matrix of BASELINE MODEL and TL model for full-dataset and reduced-dataset when trained on All-users—left and Leave-one-user-out—right.




Figure 6: Performance evaluation. Left figure shows four distribution of performance scores $((ACC + AUC)/2)$ for BASELINE and TC models for full-dataset and reduced-dataset; Right tables show Shapiro-Wilk statistics confirming normal distribution of data and Paired sample T-test with 95% confidence intervals and Cohen d effect size.

work should expand the gesture set and explore how generalisable are models to other objects.

We captured data with 10 users, however not all users participated in all sessions introducing variation of the number of recordings each user made per gesture. So when experimenting the leave-one-user-out scenario, to treat fairly all users, we chose the maximal common repetition number discarding some recordings (e.g. out of 4300 we used 2000 gesture recordings). In the future, researchers could train and experiment with the whole dataset. Future research should also be directed into few-shot learning. For example it would be interesting to see how well would transfer learning perform when trained with only a few recordings per gesture.

7 Conclusion

In this paper we demonstrates the potential of radar sensing for robust gesture recognition on deformable objects, offering a scalable alternative to sensor-embedded solutions. With a wrist-worn mmWave radar sensor, we achieve high classification accuracy (up to 90%) on a five-gesture set using a plush toy, despite the challenges posed by continuous object deformation and user variability. Importantly, we show that transfer learning significantly enhances generalizability to previously unseen users, improving recognition performance by up to 20%. These findings highlight the feasibility of radar-based interaction with deformable objects, paving the way for more spontaneous and practical digital control mechanisms in everyday scenarios.

Acknowledgments

This research was funded by the (i) Slovenian Research Agency, grant numbers N2-0354, BI-NO/25-27-007, P5-0433, IO-0035, J5-50155, J7-50096, (ii) research program CogniCom (0013103) at the University of Primorska, and (iii) by the EU EIC Pathfinder-Awareness Inside challenge “Symbiotik” project (from 1 Oct 2022 to 30 Sept 2026) under Grant no. 101071147.

References

- [1] Nuwan T. Attygalle, Luis A. Leiva, Matjaž Kljun, Christian Sandor, Alexander Plopski, Hirokazu Kato, and Klen Čopić Pucihar. 2021. No Interface, No Problem: Gesture Recognition on Physical Objects Using Radar Sensing. *Sensors* 21, 17 (2021). doi:10.3390/s21175771
- [2] Feiyu Chen, Honghao Lyu, Zhibo Pang, Junhui Zhang, Yonghong Hou, Ying Gu, Huayong Yang, and Geng Yang. 2019. WristCam: A Wearable Sensor for Hand Trajectory Gesture Recognition and Intelligent Human–Robot Interaction. *IEEE Sensors Journal* 19, 19 (2019), 8441–8451. doi:10.1109/JSEN.2018.2877978
- [3] Stefano Chioccarello, Arthur Sluÿters, Alberto Testolin, Jean Vanderdonck, and Sébastien Lambot. 2023. FORTE: Few Samples for Recognizing Hand Gestures with a Smartphone-attached Radar. *Proc. ACM Hum. Comput. Interact.* 7, EICS (2023), 1–25. doi:10.1145/3593231
- [4] Harrison Cook, Quang Vinh Nguyen, and Simeon Simoff. 2015. Enabling Finger-Gesture Interaction with Kinect. In *Proceedings of the 8th International Symposium on Visual Information Communication and Interaction* (Tokyo, AA, Japan) (VINCI '15). Association for Computing Machinery, New York, NY, USA, 152–153. doi:10.1145/2801040.2801060
- [5] Khaled Elgeneidy, Gerhard Neumann, Simon Pearson, Michael Jackson, and Niels Lohse. 2018. Contact Detection and Size Estimation Using a Modular Soft Gripper with Embedded Flex Sensors. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 498–503. doi:10.1109/IROS.2018.8593399
- [6] Barrett Ens, Aaron Quigley, Hui-Shyong Yeo, Pourang Irani, Thammathip Piumsomboon, and Mark Billinghurst. 2018. Counterpoint: Exploring Mixed-Scale Gesture Interaction for AR Applications. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI EA '18). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3170427.3188513
- [7] Tamil Selvan Gunasekaran, Ryo Hajika, Yun Suen Pai, Eiji Hayashi, and Mark Billinghurst. 2022. RaITIn: Radar-Based Identification for Tangible Interactions. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI EA '22). Association for Computing Machinery, New York, NY, USA, Article 445, 7 pages. doi:10.1145/3491101.3519808
- [8] Ryo Hajika, Tamil Selvan Gunasekaran, Chloe Dolma Si Ying Haigh, Yun Suen Pai, Eiji Hayashi, Jaime Lien, Danielle Lottridge, and Mark Billinghurst. 2024. Radar-Hand: A Wrist-Worn Radar for On-Skin Touch-Based Proprioceptive Gestures. 31, 2, Article 17 (Jan. 2024), 36 pages. doi:10.1145/3617365
- [9] Fang Hu, Peng He, Songlin Xu, Yin Li, and Cheng Zhang. 2020. FingerTrak: Continuous 3D Hand Pose Tracking by Deep Learning Hand Silhouettes Captured by Miniature Thermal Cameras on Wrist. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 2, Article 71 (June 2020), 24 pages. doi:10.1145/3397306
- [10] Gota Kakehi, Yuta Sugiura, Anusha Withana, Calista Lee, Naohisa Nagaya, Daisuke Sakamoto, Maki Sugimoto, Masahiko Inami, and Takeo Igarashi. 2011. FuwaFuwa: detecting shape deformation of soft objects using directional photoreflectivity measurement. In *ACM SIGGRAPH 2011 Emerging Technologies* (Vancouver, British Columbia, Canada) (SIGGRAPH '11). Association for Computing Machinery, New York, NY, USA, Article 5, 1 pages. doi:10.1145/2048259.2048264
- [11] Johan Kildal. 2012. Interacting with Deformable User Interfaces: Effect of Material Stiffness and Type of Deformation Gesture. In *Haptic and Audio Interaction Design*, Charlotte Magnusson, Delphine Szymczak, and Stephen Brewster (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 71–80.
- [12] Panayiotis Koutsabasis and Panagiotis Vogiatzidakis. 2019. Empirical Research in Mid-Air Interaction: A Systematic Review. *International Journal of Human–Computer Interaction* 35, 18 (2019), 1747–1768. doi:10.1080/10447318.2019.1572352
- [13] Gierad Laput, Robert Xiao, and Chris Harrison. 2016. ViBand: High-Fidelity Bio-Acoustic Sensing Using Commodity Smartwatch Accelerometers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Tokyo, Japan) (UIST '16). Association for Computing Machinery, New York, NY, USA, 321–333. doi:10.1145/2984511.2984582
- [14] Sang-Su Lee, Sohyun Kim, Bopil Jin, Eunji Choi, Boa Kim, Xu Jia, Daeeop Kim, and Kun-pyo Lee. 2010. How users manipulate deformable displays as input devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 1647–1656. doi:10.1145/1753326.1753572
- [15] Wonkyum Lee, Sungkwan Jung, Sangsik Kim, Woojin Ahn, and Sang-su Lee. 2011. Design and Implementation of Deformation Based Gesture Interaction. In *HCI International 2011 – Posters' Extended Abstracts*, Constantine Stephanidis (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 176–179.
- [16] Luis A. Leiva, Matjaž Kljun, Christian Sandor, and Klen Čopić Pucihar. 2020. The Wearable Radar: Sensing Gestures Through Fabrics. In *Proc. MobileHCI*. 1–4.
- [17] Feifei Li, Yujun Li, Baozhen Du, Hongji Xu, Hailiang Xiong, and Min Chen. 2019. A Gesture Interaction System Based on Improved Finger Feature and WE-KNN. In *Proceedings of the 4th ACM International Conference on Mathematics and Artificial Intelligence* (Chengdu, China) (ICMAI 2019). Association for Computing Machinery, New York, NY, USA, 39–43. doi:10.1145/3325730.3325759
- [18] Jaime Lien, Nicholas Gillian, M. Emre Karagozler, Patrick Amihoud, Carsten Schwesig, Erik Olson, Hakim Raja, and Ivan Poupyrev. 2016. Soli: ubiquitous gesture sensing with millimeter wave radar. *ACM Trans. Graph.* 35, 4, Article 142 (July 2016), 19 pages. doi:10.1145/2897824.2925953
- [19] Khadijeh Mahdikhani and Hossein Ebrahimnezhad. 2022. 3D object deforming and manipulating through dynamic hand gestures. *Entertainment Computing* 43 (2022), 100507. doi:10.1016/j.entcom.2022.100507
- [20] Perla Maiolino, Marco Maggiali, Giorgio Cannata, Giorgio Metta, and Lorenzo Natale. 2013. A Flexible and Robust Large Scale Capacitive Tactile System for Robots. *IEEE Sensors Journal* 13, 10 (2013), 3910–3917. doi:10.1109/JSEN.2013.2258149
- [21] Kristian Mrazek, Brian Holton, Tanner Klein, Izan Khan, Thomas Ayele, and Tauheed Khan Mohd. 2021. The TapStrap 2: Evaluating Performance of One-Handed Wearable Keyboard and Mouse. In *Proceedings of 23rd International Conference on Human-Computer Interaction - Late Breaking Papers: Multimodality, eXtended Reality, and Artificial Intelligence* (Virtual event) (HCI International 2021), Constantine Stephanidis, Masaaki Kurosu, Jessie Y. C. Chen, Gino Fragnomeni, Norbert Streitz, Shin'ichi Konomi, Helmut Degen, and Stavroula Ntoa (Eds.). Springer International Publishing, Cham, 82–95. doi:10.1007/978-3-030-90963-5_7
- [22] Katsunori Ohnishi, Atsushi Kanehira, Asako Kanezaki, and Tatsuya Harada. 2016. Recognizing Activities of Daily Living with a Wrist-Mounted Camera. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3103–3111. doi:10.1109/CVPR.2016.338
- [23] Patrick Parzer, Adwait Sharma, Anita Vogl, Jürgen Steimle, Alex Olwal, and Michael Haller. 2017. SmartSleeve: Real-time Sensing of Surface and Deformation Gestures on Flexible, Interactive Textiles, using a Hybrid Gesture Detection Pipeline. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology* (Québec City, QC, Canada) (UIST '17). Association for Computing Machinery, New York, NY, USA, 565–577. doi:10.1145/3126594.3126652
- [24] J. Rekimoto. 2001. GestureWrist and GesturePad: unobtrusive wearable interaction devices. In *Proceedings Fifth International Symposium on Wearable Computers*. 21–27. doi:10.1109/ISWC.2001.962092
- [25] Munehiko Sato, Ivan Poupyrev, and Chris Harrison. 2012. Touché: enhancing touch interaction on humans, screens, liquids, and everyday objects. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 483–492. doi:10.1145/2207676.2207743
- [26] Yukina Sato, Takashi Amesaka, Takumi Yamamoto, Hiroki Watanabe, and Yuta Sugiura. 2024. Exploring User-Defined Gestures as Input for Hearables and Recognizing Ear-Level Gestures with IMUs. *Proc. ACM Hum.-Comput. Interact.* 8, MHCI, Article 258 (Sept. 2024), 23 pages. doi:10.1145/3676503
- [27] Alexander Schmitz, Marco Maggiali, Marco Randazzo, Lorenzo Natale, and Giorgio Metta. 2008. A prototype fingertip with high spatial resolution pressure sensing for the robot iCub. In *Humanoids 2008 - 8th IEEE-RAS International Conference on Humanoid Robots*. 423–428. doi:10.1109/ICHR.2008.4755988
- [28] Alexander Schmitz, Perla Maiolino, Marco Maggiali, Lorenzo Natale, Giorgio Cannata, and Giorgio Metta. 2011. Methods and Technologies for the Implementation of Large-Scale Robot Tactile Sensors. *IEEE Transactions on Robotics* 27, 3 (2011), 389–400. doi:10.1109/TRO.2011.2132930
- [29] Adwait Sharma, Alexander Ivanov, Frances Lai, Tovi Grossman, and Stephanie Santosa. 2024. GraspUI: Seamlessly Integrating Object-Centric Gestures within the Seven Phases of Grasping. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) (DIS '24). Association for Computing Machinery, New York, NY, USA, 1275–1289. doi:10.1145/3643834.3661551
- [30] Adwait Sharma, Joan Sol Roo, and Jürgen Steimle. 2019. *Grasping Microgestures: Eliciting Single-Hand Microgestures for Handheld Objects*. 1–13.
- [31] Ewa Siedlecka and Thomas F. Denson. 2019. Experimental Methods for Inducing Basic Emotions: A Qualitative Review. *Emotion Review* 11, 1 (2019), 87–97. doi:10.1177/1754073917749016 arXiv:https://doi.org/10.1177/1754073917749016
- [32] Arthur Sluÿters, Sébastien Lambot, and Jean Vanderdonck. 2022. Hand Gesture Recognition for an Off-the-Shelf Radar by Electromagnetic Modeling and Inversion. In *27th International Conference on Intelligent User Interfaces* (Helsinki, Finland) (IUI '22). ACM, New York, NY, USA, 506–522. doi:10.1145/3490099.3511107

- [33] Arthur Sluÿters, Sébastien Lambot, Jean Vanderdonckt, and Radu-Daniel Vatavu. 2023. RadarSense: Accurate Recognition of Mid-Air Hand Gestures with Radar Sensing and Few Training Examples. *ACM Trans. Interact. Intell. Syst.* (mar 2023). doi:10.1145/3589645 Just Accepted.
- [34] Arthur Sluÿters, Sébastien Lambot, Jean Vanderdonckt, and Santiago Villarreal-Narvaez. 2024. Analysis of User-Defined Radar-Based Hand Gestures Sensed Through Multiple Materials. *IEEE Access* 12 (2024), 27895–27917. doi:10.1109/ACCESS.2024.3366667
- [35] Arthur Sluÿters, Mehdi Ousmer, Paolo Roselli, and Jean Vanderdonckt. 2022. QuantumLeap, a Framework for Engineering Gestural User Interfaces based on the Leap Motion Controller. *Proc. ACM Human-Computer Interaction* 6, EICS, Article 161 (jun 2022), 47 pages. doi:10.1145/3532211
- [36] Arthur Sluÿters, Quentin Sellier, Jean Vanderdonckt, Vik Parthiban, and Pattie Maes. 2023. Consistent, Continuous, and Customizable Mid-Air Gesture Interaction for Browsing Multimedia Objects on Large Displays. *International Journal of Human-Computer Interaction* 39, 12 (2023), 2492–2523. doi:10.1080/10447318.2022.2078464 arXiv:https://doi.org/10.1080/10447318.2022.2078464
- [37] Anne-Sophie Tribot, Nathalie Blanc, Thierry Brassac, François Guilhaumon, Nicolas Casajus, and Nicolas Mouquet. 2024. What makes a teddy bear comforting? A participatory study reveals the prevalence of sensory characteristics and emotional bonds in the perception of comforting teddy bears. *The Journal of Positive Psychology* 19, 2 (2024), 379–392. doi:10.1080/17439760.2023.2170273 arXiv:https://doi.org/10.1080/17439760.2023.2170273
- [38] Ayumu Tsuboi, Mamoru Hirota, Junki Sato, Masayuki Yokoyama, and Masao Yanagisawa. 2017. A proposal for wearable controller device and finger gesture recognition using surface electromyography. In *Proceedings of SIGGRAPH Asia 2017 Posters* (Bangkok, Thailand) (SA '17). Association for Computing Machinery, New York, NY, USA, Article 9, 2 pages. doi:10.1145/3145690.3145731
- [39] Klen Čopić Pucihar, Christian Sandor, Matjaž Kljun, Wolfgang Huerst, Alexander Plopski, Takafumi Taketomi, Hirokazu Kato, and Luis A. Leiva. 2019. The Missing Interface: Micro-Gestures on Augmented Objects. In *Proceedings of Extended Abstracts on Human Factors in Computing Systems (CHI EA)*. 1–6.
- [40] Santiago Villarreal-Narvaez, Arthur Sluÿters, Jean Vanderdonckt, and Efreem Mbaki Luzayisu. 2022. Theoretically-Defined vs. User-Defined Squeeze Gestures. *Proc. ACM Hum.-Comput. Interact.* 6, ISS, Article 559 (Nov. 2022), 30 pages. doi:10.1145/3567805
- [41] Cong Wang, Deepak Vungarala, Kevin Navarro, Neel Adwani, and Tao Han. 2023. The Pinch Sensor: An Input Device for In-Hand Manipulation with the Index Finger and Thumb. In *2023 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. 822–827. doi:10.1109/AIM46323.2023.10196268
- [42] Saiwen Wang, Jie Song, Jaime Lien, Ivan Poupyrev, and Otmár Hilliges. 2016. Interacting with Soli: Exploring Fine-Grained Dynamic Gesture Recognition in the Radio-Frequency Spectrum. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (Tokyo, Japan) (UIST '16)*. Association for Computing Machinery, New York, NY, USA, Article 189, 4 pages. doi:10.1145/2984511.2984565
- [43] Keisuke Watanabe, Ryosuke Yamamura, and Yasuaki Kakehi. 2021. foamIn: A Deformable Sensor for Multimodal Inputs Based on Conductive Foam with a Single Wire. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '21). Association for Computing Machinery, New York, NY, USA, Article 189, 4 pages. doi:10.1145/3411763.3451547
- [44] HS Yeo, A Quigley Interactions, and Undefined. 2017. Radar Sensing in Human-computer Interaction. *Interactions* 25, 1 (2017), 70–73.
- [45] Renyuan Zhang and Siyang Cao. 2018. Real-time human motion behavior detection via CNN using mmWave radar. *IEEE Sens. Lett.* 3, 2 (2018), 1–4.
- [46] Klen Čopić Pucihar, Nuwan T. Attygalle, Matjaž Kljun, Christian Sandor, and Luis A. Leiva. 2022. Solids on Soli: Millimetre-Wave Radar Sensing through Materials. *Proc. ACM Human-Computer Interaction* 6, EICS, Article 156 (jun 2022), 21 pages. doi:10.1145/3532212

A Appendix

In Table 2 and Table 3 are results of Leave-one-user-out experimentation.

Model	Dataset	Excluded	ACC	AUC	Precision	Recall	F1
Baseline	Reduced-d.	User 1	0.65	0.94	0.68	0.65	0.65
Baseline	Reduced-d.	User 2	0.36	0.80	0.54	0.36	0.32
Baseline	Reduced-d.	User 3	0.23	0.52	0.09	0.23	0.13
Baseline	Reduced-d.	User 4	0.35	0.73	0.50	0.35	0.35
Baseline	Reduced-d.	User 5	0.33	0.68	0.42	0.33	0.30
Baseline	Reduced-d.	User 6	0.40	0.79	0.65	0.40	0.34
Baseline	Reduced-d.	User 7	0.92	0.99	0.93	0.92	0.92
Baseline	Reduced-d.	User 8	0.32	0.70	0.33	0.32	0.20
Baseline	Reduced-d.	User 9	0.73	0.92	0.75	0.73	0.73
Baseline	Reduced-d.	User 10	0.45	0.80	0.46	0.45	0.40
Median			0.38	0.80	0.52	0.38	0.35
Mean			0.47	0.79	0.54	0.47	0.43
STD			0.22	0.14	0.23	0.22	0.25
Baseline	Full-d.	User 1	0.76	0.94	0.79	0.76	0.76
Baseline	Full-d.	User 2	0.81	0.98	0.83	0.81	0.80
Baseline	Full-d.	User 3	0.75	0.95	0.80	0.75	0.75
Baseline	Full-d.	User 4	0.62	0.90	0.62	0.62	0.59
Baseline	Full-d.	User 5	0.85	0.98	0.86	0.85	0.85
Baseline	Full-d.	User 6	0.67	0.93	0.68	0.67	0.66
Baseline	Full-d.	User 7	0.93	0.99	0.93	0.93	0.93
Baseline	Full-d.	User 8	0.86	0.99	0.88	0.86	0.85
Baseline	Full-d.	User 9	0.60	0.88	0.64	0.59	0.59
Baseline	Full-d.	User 10	0.75	0.95	0.76	0.75	0.74
Median			0.76	0.95	0.80	0.76	0.76
Mean			0.76	0.95	0.78	0.76	0.75
STD			0.11	0.04	0.10	0.11	0.11

Table 2: Leave-one-user-out performance evaluation of BASELINE MODEL and TL model for full-dataset and reduced-dataset.

Model	Dataset	Excluded	ACC	AUC	Precision	Recall	F1
TL	Reduced-d.	User 1	0.69	0.92	0.72	0.69	0.70
TL	Reduced-d.	User 2	0.56	0.93	0.64	0.56	0.52
TL	Reduced-d.	User 3	0.73	0.92	0.75	0.73	0.73
TL	Reduced-d.	User 4	0.49	0.79	0.46	0.49	0.46
TL	Reduced-d.	User 5	0.77	0.94	0.78	0.77	0.77
TL	Reduced-d.	User 6	0.67	0.92	0.68	0.67	0.66
TL	Reduced-d.	User 7	0.91	0.99	0.91	0.91	0.91
TL	Reduced-d.	User 8	0.89	0.99	0.90	0.89	0.89
TL	Reduced-d.	User 9	0.73	0.92	0.77	0.73	0.74
TL	Reduced-d.	User 10	0.73	0.94	0.78	0.73	0.72
Median			0.73	0.93	0.76	0.73	0.73
Mean			0.72	0.93	0.74	0.72	0.71
STD			0.13	0.06	0.13	0.13	0.14
TL	Full-d.	User 1	0.76	0.95	0.77	0.76	0.76
TL	Full-d.	User 2	0.76	0.97	0.80	0.77	0.76
TL	Full-d.	User 3	0.80	0.96	0.81	0.80	0.80
TL	Full-d.	User 4	0.81	0.97	0.81	0.81	0.80
TL	Full-d.	User 5	0.88	0.99	0.80	0.88	0.83
TL	Full-d.	User 6	0.92	1.00	0.92	0.92	0.92
TL	Full-d.	User 7	0.91	0.99	0.92	0.91	0.91
TL	Full-d.	User 8	0.87	0.99	0.88	0.87	0.86
TL	Full-d.	User 9	0.92	1.00	0.92	0.92	0.92
TL	Full-d.	User 10	0.81	0.96	0.84	0.81	0.81
Median			0.84	0.98	0.83	0.84	0.82
Mean			0.84	0.98	0.85	0.85	0.84
STD			0.06	0.02	0.06	0.06	0.06

Table 3: Leave-one-user-out performance evaluation of BASELINE MODEL and TL model for full-dataset and reduced-dataset.