



# Activity classification using unsupervised domain transfer from body worn sensors

Chaitra Hegde <sup>a</sup>, Gezheng Wen <sup>b</sup>, Layne C. Price <sup>b,\*</sup>

<sup>a</sup> School of Electrical and Computer Engineering, Georgia Institute of Technology, 777 Atlantic Drive NW, Atlanta, GA 30332, USA

<sup>b</sup> Amazon Services Inc., 2121 7th Ave, Seattle, WA 98121, USA

## ARTICLE INFO

### Keywords:

Domain adaptation  
Wearable sensing  
Activity recognition  
Unsupervised learning

## ABSTRACT

Activity classification has become a vital feature of wearable health tracking devices. As innovation in this field grows, wearable devices worn on different parts of the body are emerging. To perform activity classification on a new body location, labeled data corresponding to the new locations are generally required, but this is expensive to acquire. In this work, we present an innovative method to leverage an existing activity classifier, trained on Inertial Measurement Unit (IMU) data from a reference body location (the source domain), in order to perform activity classification on a new body location (the target domain) in an unsupervised way, i.e. without the need for classification labels at the new location. Specifically, given an IMU embedding model trained to perform activity classification at the source domain, we train an embedding model to perform activity classification at the target domain by replicating the embeddings at the source domain. This is achieved using simultaneous IMU measurements at the source and target domains. The replicated embeddings at the target domain are used by a classification model that has previously been trained on the source domain to perform activity classification at the target domain. We have evaluated the proposed methods on three activity classification datasets PAMAP2, MHealth, and Opportunity, yielding high F1 scores of 67.19%, 70.40% and 68.34%, respectively when the source domain is the wrist and the target domain is the torso.

## 1. Introduction

Body-worn measurement devices have become ubiquitous in both consumer and health contexts, as they can provide real-time monitoring and inference on a person's activities, as well as their general health and wellness. Human activity recognition has been shown to be valuable in various fields, such as healthcare (Badar ud din Tahir, Jalal, & Kim, 2021; Zhou et al., 2020), patient monitoring (Gjoreski, Gjoreski, Luštrek, & Gams, 2016; Gjoreski, Lustrek, & Gams, 2011; Li, Shrestha, Heidari, Le Kernec, & Fioranelli, 2020; Syed, Sierra-Sosa, Kumar, & Elmaghraby, 2022), gait analysis (Caramia et al., 2018; Zhao, Wang, Qiu, Shen, & Wang, 2017), surveillance and safety applications (Al-Nawashi, Al-Hazaimah, & Saraei, 2017; Gowsikhaa, Abirami, et al., 2012; Taha, Zayed, Khalifa, & El-Horbaty, 2015), and gesture recognition (Kim, Cho, Lee, & Jung, 2019; Nunez, Cabido, Pantrigo, Montemayor, & Velez, 2018). Most of these human activity recognition methods use machine learning or statistical methods based on accelerometers and/or inertial measurement unit (IMU) sensors present onboard the device, either independently or in conjunction with other sensors, such as photoplethysmogram (Alessandrini, Biagetti, Crippa, Falaschetti, & Turchetti, 2021; Mekruksavanich, Jantawong, Hnoohom, & Jitpattanakul, 2022).

\* Corresponding author.


E-mail address: [prilayne@amazon.com](mailto:prilayne@amazon.com) (L.C. Price).

<https://doi.org/10.1016/j.smhl.2023.100431>

Received 20 April 2023; Received in revised form 28 August 2023; Accepted 25 September 2023

Available online 11 October 2023

2352-6483/© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).



**Fig. 1.** An example of source domain (e.g., wrist) and target domain (e.g., torso) body locations. This work uses a model that is trained to perform activity classification on the source domain to do the same on the target domain. It is assumed that a well performing activity classifier  $M_S$  has been trained on the source domain data and that a small amount of unlabeled data is available for the target domain with simultaneous readings from the source domain.  $M_S$  is retrained using this unlabeled data to obtain  $M_T$  which is used for activity classification on the target domain.

Wrist-worn measurements devices such as smart watches and fashionable accessories can easily host an IMU, and are hence the most common consumer-grade activity measurement tool. There are plenty of good quality labeled data from the wrist, on which numerous advanced activity classification algorithms have been trained. There is growing interest in activity recognition from body-worn measurement devices that are not placed on the wrist, e.g., devices placed on the ankle, torso, or waist (Dobkin, Xu, Batalin, Thomas, & Kaiser, 2011; Storm, Buckley, & Mazzà, 2016). IMU devices placed on the chest can be used to perform various functions like seismocardiography parameters estimation, voice analysis, posture analysis, amongst many others (Rahmani, Berkvens, & Weyn, 2021). IMUs placed at the ankle can be used for analysis of ankle joint movement during running (Kim et al., 2021). IMUs are increasingly being placed in earbuds which are being used for various uses like cough detection (Lamichhane et al., 2022) and eating detection (Bin Morshed, Haresamudram, Bandaru, Abowd, & Ploetz, 2022). Due to activity recognition's popularity, especially for exercise tracking, this feature has become a minimum expectation in most new wearable devices. Having the capability to perform activity recognition using IMU sensors that are already present on these body parts will be beneficial rather than having to wear a secondary device for this purpose. There is growing interest in body-worn measurement devices that are not placed on the wrist, e.g., devices placed on the ankle, torso, or waist (Dobkin et al., 2011; Storm et al., 2016), but the amount of labeled data from these body locations are scarce. In this paper we propose a new method to use an existing activity classifier that has been trained on IMU data obtained in a source domain (e.g. the wrist) to perform activity classification predictions on a target domain (e.g. the torso) in an unsupervised way. Fig. 1 displays example traces of unlabeled, simultaneously collected IMU data from the source and target domains.

Our method falls under the umbrella of unsupervised domain transfer methods. It does not require labeled data at the target body site, and only relies on simultaneous IMU measurements at the source and target body sites. A large amount of such data can be easily obtained in various day-to-day scenarios from a large number of subjects, thus capturing a wide variety of complex activities. Additionally, our method can be generalized to any source and target combinations. A unique characteristic of our method is that an existing activity classifier trained on the source body location can be applied at the target body location without making changes to the architecture of the classifier. Thus, advantages of the trained classifier can be retained, and applied directly to the target body site, which would significantly save cost and modeling effort. This is especially useful in commercial settings where well performing activity classification models already exist for wrist IMU data, and are efficient enough to run on edge devices.

## 2. Related work

Multiple measurement modalities have been used for activity recognition, including body worn sensors like inertial measurement units (Hassan, Uddin, Mohamed, & Almogren, 2018; Luwe, Lee, & Lim, 2022; Zhu, Chen, & Soh, 2019) and accelerometers (Sukor, Zakaria, & Rahim, 2018), off-body sensors like cameras (Almaadeed, Elharrouss, Al-Maadeed, Bouridane, & Beghdadi, 2019; Gleason et al., 2018) and infrared sensors (Wren & Tapia, 2006), and a combination of sensors as used in Ogris, Stiefmeier, Lukowicz, and Troster (2008) and Ordóñez and Roggen (2016).

A large number of algorithms in literature use supervised machine learning methods to perform activity classification using data from these sensors, such as support vector machines (Hassan et al., 2018; Köping, Shirahama, & Grzegorzec, 2018; Sukor et al., 2018), convolutional neural networks (CNNs) (Khan & Ahmad, 2021; Moya Rueda, Grzeszick, Fink, Feldhorst, & Ten Hompel, 2018; Qi et al., 2019), long-short term memory networks (LSTMs) (Saini, Kumar, Roy, & Dogra, 2018), attention mechanisms (Khan & Ahmad, 2021; Ma, Li, Zhang, Gao, & Lu, 2019), and a combination of CNNs and recurrent neural networks (Luwe et al., 2022; Ma et al., 2019).

Supervised machine learning approaches generally require large amounts of labeled data. The quality of the labeled data has a direct impact on the performance of the supervised model. However, collecting and labeling high-quality data from human subjects is expensive and time-consuming and thus not scalable to real-world problems.

Many works in the literature have identified this shortcoming and have proposed methods to circumvent this problem to the extent possible. Active learning techniques have been utilized to reduce labeling effort by partially automating the labeling

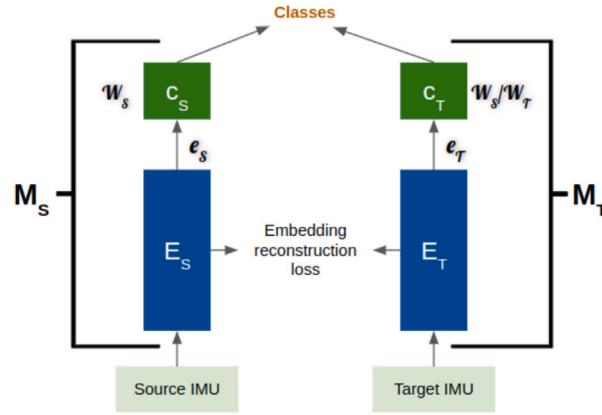


Fig. 2. Diagram showing the unsupervised domain adaptation approach. The source-domain model  $M_S$  consists of the embedding extractor module  $E_S$  and the classifier module  $C_S$ . Similarly the target-domain model  $M_T$  consists of the embedding extractor module  $E_T$  and classifier module  $C_T$ . The embedding extractor modules  $E_S$  and  $E_T$  generate embeddings  $e_S$  and  $e_T$  from the input IMU signals. These embeddings are used by the classifier modules  $C_S$  and  $C_T$  to output the activity classes for the source and target domains respectively. An embedding reconstruction loss is used to train  $E_T$  with IMU signals simultaneously recorded at the source and target body locations such that  $e_T$  is similar to  $e_S$ . The weights  $W_S$  of classifier  $C_S$  can be applied to classifier  $C_T$  when no target domain labels are available. In the case where target domain labels are available,  $C_T$  can be fine-tuned using the labels to have weights  $W_T$ .

process (Bota, Silva, Folgado, & Gamboa, 2019; Stikic, Van Laerhoven, & Schiele, 2008). Semi-supervised learning methods use a small amount of labeled data and a large amount of unlabeled data to train models (Chen et al., 2020; Guan, Yuan, Lee, Gavrilov, & Lee, 2007; Stikic et al., 2008; Zhu et al., 2019). Weakly supervised learning techniques have been used to allow the use of small amounts of sparsely labeled data for training (Stikic, Larlus, Ebert, & Schiele, 2011). Ensembles of models have also been trained on small quantities of sensor data (Chen et al., 2020). Unsupervised learning methods do not rely on labeled data and discover structures in unlabeled data (Kwon, Kang, & Bae, 2014; Wyatt, Philipose, & Choudhury, 2005).

Sometimes, a large amount of labeled data are available for activity classification at a particular body location (the source domain) and a small amount of labeled data is available at the target body location where an activity classifier is desired. Transfer learning approaches have been commonly used in such cases (Chen, Wang, Huang and Yu, 2019; Wang, Chen, Hu, Peng and Yu, 2018; Wang, Zheng, Chen and Huang, 2018), and other scenarios, such as between different user populations (Li, Kang, Tan, & Shull, 2022), body locations (Chen, Wang et al., 2019; Wang, Chen et al., 2018), application domains and sensor modalities (Morales & Roggen, 2016). Stratified transfer learning methods have been popular in such applications (Chen, Wang et al., 2019; Wang, Chen et al., 2018). Another approach that has been used is to add a new layer to an existing pre-trained model in addition to retraining the model (Wang, Zheng et al., 2018).

Unsupervised domain adaptation is an approach used when no labeled data is available at the target body site. Faridee, Chakma, Misra, and Roy (2022) used a spatial transformer to learn affine transformations that transform the target IMU to have discriminative features similar to those for the source domain IMU. Chen, Miao et al. (2019) proposed a more complex model with encoder, generator, decoder and predictor modules to convert raw IMU signals into domain invariant hidden representations. Mu, Gu, Guo, and Lo (2020) used domain adversarial neural networks (DANN) (Ganin et al., 2016) and extended multi-source DANN (Pei, Cao, Long, & Wang, 2018; Zhao et al., 2018) to perform unsupervised domain adaptation for gait event detection and abnormal gait pattern recognition, in which DANN projects IMU data from the source and target domains into a common subspace using a feature extractor, domain classifier and a task specific classifier.

Our method is compared with Faridee et al. (2022) due to the use of similar datasets. Our work offers comparable or better performance than (Faridee et al., 2022).

### 3. Unsupervised domain transfer approach

The schematic of our proposed unsupervised domain transfer approach is shown in Fig. 2. Our approach consists of two models, the source-domain model  $M_S$  and the target-domain model  $M_T$ . The source-domain model,  $M_S$ , can either be an existing pre-trained model or it can be trained in a supervised way using a source-domain dataset  $D_S$  containing a collection of IMU or accelerometer signals  $\alpha_S(t)$  and labels  $l_S(t)$  and can be defined as:

$$D_S = \{[\alpha_S(t), l_S(t)]\} \quad (1)$$

$M_S$  can be decomposed into an embedding extractor module  $E_S$  and a classification module  $C_S$ . The embedding extractor module  $E_S$  generates embeddings  $e_S$  from the input IMU signals while the classification module uses these embeddings to output the final classified labels.  $M_S$  can be defined as:

$$M_S^\theta(\alpha_S(t)) = \hat{l}_S(t), \quad (2)$$

where  $\theta$  are the real parameters of this model and  $\hat{l}_S$  are the model's estimated output labels.  $M_S$  should have the following structure:

$$M_S^\theta(\alpha_S) = \phi(W_S e_S^\theta(\alpha_S)), \quad (3)$$

where  $\phi$  is a multi-dimensional logistic function,  $W_S$  is a real-valued weight matrix, and  $e_S^\theta \in \mathbb{R}^d$  is the  $d$ -dimensional embedding for the signal  $\alpha_S$  produced by the embedding extractor module ( $E_S$ ) of the model.

Similarly, the target-domain model  $M_T$  can be decomposed into an embedding extractor module  $E_T$  and a classification module  $C_T$ . The embedding extractor module  $E_T$  generates embeddings  $e_T$  from the input IMU signals while the classification module uses these embeddings to output the final classified labels.  $M_T$  can be defined as:

$$M_T^\psi = \phi(W_T e_T^\psi(\alpha_T)), \quad (4)$$

where  $W_T$  is a real-valued weight matrix,  $e_T \in \mathbb{R}^d$ , and  $\psi$  are real-valued parameters of the model.

The goal of our procedures is to train  $E_T$  such that the produced embeddings  $e_T$  are as similar to  $e_S$  as possible for each activity class. For this, we use a dataset  $D_{ST}$  that has simultaneous IMU or accelerometer signals measured at the source body location and the target body location, but does not necessarily contain activity labels. These time series signals at the source and target body locations are  $\alpha_S(t)$  and  $\alpha_T(t)$ , respectively. This dataset  $D_{ST}$  can be defined as:

$$D_{ST} = \{[\alpha_S(t), \alpha_T(t)]\}. \quad (5)$$

$E_S$  provides a set of features,  $e_S^\theta(\alpha_S)$ , that are relevant for predicting the labels  $l_S$ . In this work, we further make the critical assumption that one or more of the features in  $e_S$  are relevant for making predictions for the target-domain labels  $l_T$ . In other words, we assume that most dimensions in the source domain embeddings  $e_S$  can be reconstructed from the target domain signals  $\alpha_T(t)$ . It is important to note that only certain activity classes would have this property. For example, if the source domain is the wrist and the target domain is the torso, activities such as walking and jumping contain discriminating information that is present in both domains, as both body sites physically participate in such activity. In contrast, activities like typing or clapping have relevant information only in the wrist but not in the torso; typing and clapping are uninformative in the target domain. Here, in the scope of this paper, we restrict our attention to activity labels that can be mutually predicted from both source and target domain measurements. Therefore, our approach is to obtain target-domain embeddings  $e_T^\psi$  by minimizing a loss function  $L$  that has a global minimum where  $e_T(\alpha_T) = e_S(\alpha_S)$  for  $[\alpha_T, \alpha_S] \in D_{ST}$ :

$$\hat{\psi} = \underset{\psi}{\operatorname{argmin}} L[e_S^\theta(\alpha_S), e_T^\psi(\alpha_T)] \quad (6)$$

Choices of the loss function  $L$  can be varied; e.g., in Section 6 we consider mean-absolute (MAE), mean-squared (MSE), and mean-squared-logarithmic error (MSLE), as well as cosine similarity. Regularizations such as  $L_1$  and  $L_2$  regularizations can be also applied to promote the reconstruction of subsets of the source-domain embeddings that are most relevant to making predictions in the target domain.

Finally, we can obtain an unsupervised activity classifier at the target domain that predicts the source domain labels  $l_S$  by setting  $W_T = W_S$  in Eq. (4). Alternatively, if target-domain labels  $l_{ST}$  are available in  $D_{ST}$ ,  $W_T$  can be obtained by a second supervised fine-tuning step with a classification-relevant loss function, similarly to how  $M_S$  is trained.

#### 4. Implementation

The implementation of the approach discussed in Section 3 is explained here. The models  $M_S$  and  $M_T$  in Eqs. (3) and (4) are comprised of two components: the embedding extractor and the classifier. For both models  $M_S$  and  $M_T$ , we used the state-of-the-art convolution-LSTM architecture proposed by Ordóñez and Roggen (2016), as shown in Fig. 3. The key advantages of the chosen architecture are its relatively high performance and ability to disambiguate closely related activities while being a relatively light-weight model. The model consists of 4 convolutional layers with 64 feature maps and ReLU activation, 2 LSTM layers with 128 cells, and a dense layer with softmax activation as the last layer. The input to the model has the shape ( $C \times W$ ) where  $C$  is the number of channels from the sensor and  $W$  is the length of the window of sensor data which we have considered to be 5 s. In this work,  $C$  is 3 for triaxial accelerometers and 9 for IMUs which collect data along all three axes. The layers up to the final LSTM layer are considered to make up the embedding extractor module ( $E_S$  and  $E_T$ ), while the rest of the layers comprise the classifier module ( $C_S$  and  $C_T$ ). Note that the architecture and hyperparameter settings for  $M_T$  and  $M_S$  are the same.

For each of the publicly available datasets described later in Section 5, we followed the methods in Ordóñez and Roggen (2016) and trained the corresponding reference source-domain model  $M_S$  in a supervised manner using cross-entropy loss function and mini-batch gradient descent with RMSprop. Note that Ordóñez and Roggen (2016) used data from multiple on-body sensors to train their model (i.e., all the sensors available in the datasets), while we only used data from the wrist sensors. For example, for the Opportunity dataset with IMU readings from 7 body locations, Ordóñez and Roggen (2016) used all 63 channels of data for training  $M_S$  whereas we used only the 9 channels from the right wrist.

The embeddings for the target-domain model  $e_T$  are found by minimizing a reconstruction loss function, as shown in Eq. (6). In Section 6 we compare results from various loss functions, including MSE, MAE, MSLE, and cosine similarity, along with  $L_1$  and  $L_2$  regularizations. We initialized the parameters of the target-domain model with those from the source-domain model, and used mini-batch gradient descent with RMPprop for optimization.




Fig. 3. Architecture diagram used for  $M_S$  and  $M_T$  obtained from Ordóñez and Roggen (2016). The layers until the final LSTM layer are considered to make up the embedding extractor modules  $E_S$  and  $E_T$ , while the rest of the layers comprise the classifier modules  $C_S$  and  $C_T$ . The source and target domain models,  $M_S$  and  $M_T$ , have the same architecture and hyperparameters.

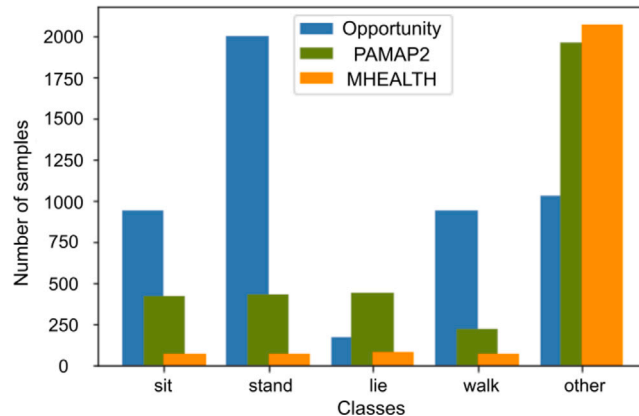


Fig. 4. Class distribution across the five overlapping labels (*sit*, *stand*, *lie down*, *walk* and *other*) for Opportunity, PAMAP2 and MHEALTH datasets. Opportunity dataset has a higher number of samples than MHEALTH and PAMAP2 datasets. PAMAP2 and MHEALTH datasets have a much larger proportion of their samples in the *other* class as compared to the rest of the classes. Each dataset is imbalanced.

## 5. Data

We used three publicly available activity classification datasets for evaluating the proposed method: Opportunity (Roggen et al., 2010), PAMAP2 (Reiss & Stricker, 2012) and MHEALTH (Banos et al., 2014). Given the relatively higher quantity of data in the Opportunity dataset, we present our results on the Opportunity dataset in details, and use the PAMAP2 and MHEALTH datasets to further validate the observations. As all three datasets have measurements at the wrist and torso/chest, we chose to use the wrist as the source domain and the torso as the target domain for our analysis.

The activity classes that overlap in all three datasets are *sitting*, *standing*, *lying down*, *walking* and *other*. For a major part of our analysis, we used these five overlapping labels so as to maintain uniformity across the datasets. The *other* class contains the remaining activities in the datasets. For the Opportunity dataset this includes activities like cleaning, relaxing, stirring, opening and closing doors. For the PAMAP2 dataset, this includes watching TV, computer work, running, car driving, vacuum cleaning, ironing, folding laundry and playing soccer and for MHEALTH dataset this includes running, ascend and descend stairs, jumping, cycling, bending forward at waist, frontal elevation of arms, knees bending and jogging. Thus the *other* class is not consistent across the three datasets. The class distribution of these five labels for each of the datasets are shown in Fig. 4. We describe each dataset below.

- The Opportunity dataset (Roggen et al., 2010) were collected in a sensor-rich environment, with on-body sensors as well as external sensors that were placed on objects that study participants interacted with. A custom motion jacket was worn by each participant that was fitted with IMUs that measured seven body locations: right lower arm, left lower arm, right upper arm, left upper arm, back, right shoe, and left shoe. We used the IMU readings from the right lower arm and the back in our analysis. This dataset has four subjects that performed two sets of experiments: activities of daily living in a routine morning, and repeated scripted sessions following a protocol. Activity classifications labels are available at multiple granularity levels: locomotion labels, high-level activities, mid-level gestures, and low-level left/right hand actions. We used only the locomotion activities for our analysis, which consists of the following activities: *sit*, *stand*, *walk*, *lie down*, and *other*.
- The PAMAP2 dataset (Reiss & Stricker, 2012) consists of 9 subjects, 8 males and 1 female in the approximate age range of 24 years to 30 years, wearing 3 IMU sensors over the dominant arm's wrist, chest, and dominant leg's ankle. Subjects in this dataset engaged in 18 day-to-day and sporting activities, including lying down, standing, walking, running, computer work, car driving, and others. For our analysis, we only focus on using the IMU readings from the dominant wrist and the chest.
- The MHEALTH dataset (Banos et al., 2014) has 10 subjects wearing an accelerometer and ECG sensor on the chest, and IMUs on the left ankle and right lower arm. The subjects performed 12 sets of physical activities, which included standing still, sitting, lying down, walking, and others. For our analysis, we used the accelerometer sensors on the chest and right wrist.

**Table 1**

Activity classification performance metrics of (1) supervised source model on the source domain, i.e.  $M_S$  on  $D_S$ , (2) supervised source model on the target domain, i.e.  $M_S$  on  $D_{ST}$  and (3) unsupervised target model on the target domain, i.e.  $M_T$  on  $D_{ST}$  for the three datasets. Training on  $D_{ST}$  as part of the modeling of  $M_T$  helps achieve better performance than  $M_S$  applied to  $D_{ST}$ , implying that the embeddings carry useful overlapping information between the source and target domains that the model  $M_T$  is able to pick up on.

Metrics	Opportunity			PAMAP2			MHEALTH		
	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$
Accuracy	76.13	24.01	65.13	94.47	61.00	77.26	92.64	80.16	91.20
Precision	73.28	14.01	66.50	94.24	38.02	73.98	73.10	19.63	69.20
Recall	79.60	26.2	67.89	92.85	41.29	67.15	88.33	24.96	67.50
F1 score	76.31	18.26	<b>67.19</b>	93.54	39.59	<b>70.40</b>	80.00	21.98	<b>68.34</b>

Note that we focused on using the accelerometer readings instead of IMU readings for this dataset since the chest was fitted with an accelerometer instead of an IMU.

The sampling frequency of the Opportunity dataset was 30 Hz, and to achieve the same sampling frequency across the three datasets, we downsampled the PAMAP2 and MHEALTH dataset to 30 Hz. Windows of data were then created such that each window had 100 samples, i.e., approximately 3 s of non-overlapping data. All datasets were made to have the same units for accelerometer ( $m/s^2$ ), gyroscope ( $deg/s$ ) and magnetometer ( $\mu T$ ) readings.

For each dataset, 30% of the dataset were randomly chosen for training the source-domain model  $M_S$ , where the wrist was considered to be the source-domain. An additional 50% of the data were randomly chosen for the unsupervised training to obtain the target-domain model  $M_T$  at the chest/torso, and the remaining 20% was used for testing.

## 6. Experiments & results

### 6.1. Performance metrics

To assess the performance of the proposed method, we computed accuracy, precision, recall and F1 scores on the provided activity classes. The precision, recall and F1 score for each of the activity classes were calculated separately by considering the classification task as a one-vs-rest problem. The scores were then averaged across the classes to obtain the final precision, recall and F1 scores. For each of the three dataset, we evaluated the performance metrics with three experiments:

- *Supervised source model on the source domain.*— The source domain model  $M_S$  trained on the source-domain data  $D_S$ . This is standard supervised training, which is used to gauge the overall performance potential of the model architecture on each of the three datasets.
- *Supervised source model on the target domain.*— The source domain model  $M_S$  trained on the source-domain data  $D_S$ , but evaluated without any fine-tuning on the target-domain data  $D_{ST}$ . This assesses how well  $M_S$  performs on  $D_{ST}$ , and provides a untrained baseline to compare the performance of trained  $M_T$  with.
- *Unsupervised target model on the target domain.*— The target-domain model  $M_T$  trained on the simultaneous source-domain and target-domain data  $D_{ST}$  using our unsupervised domain transfer method in Section 3.

Table 1 shows the accuracy, precision, recall and F1 score for each of these three models for all the three datasets. The supervised training of  $M_S$  on  $D_S$  yields the highest scores for the PAMAP2 dataset. Lower performance on the MHEALTH dataset could be partially attributed to the smaller volume of training data and fewer channels of data being used (as stated in Section 5 we only used three accelerometer channels for MHEALTH dataset as opposed to nine IMU channels for the other two datasets). As demonstrated in Guan and Plötz (2017), activities in the Opportunity dataset were less structured and closer to real-life scenarios, hence is the most difficult dataset among the three, which resulted in the lowest performance of  $M_S$  on  $D_S$ . The same performance trend across the three datasets is also observed in the other two experiments,  $M_S$  on  $D_{ST}$  and  $M_T$  on  $D_{ST}$ .

When comparing the model performance on the target domain, the target domain model  $M_T$  trained on  $D_{ST}$  performs significantly better in all the three datasets than the source domain model being directly applied to the target domain i.e.,  $M_S$  on  $D_T$ . This shows that the embedding extractor  $E_T$ , when trained on  $D_{ST}$ , is able to learn similar discriminative information as in the source domain embedding  $e_S$  for constructing the target domain embeddings  $e_T$ . Thus, fine-tuning an existing activity classification model, as part of our unsupervised method, could be potentially sufficient for performing the classification task at a new body location. However, as the performance of  $M_T$  on  $D_{ST}$  across the three datasets are in the same order as of  $M_S$  on  $D_S$ , the potential performance increase from such fine-tuning could be partially limited by the performance of the original supervised model from the source location. The outperformance of  $M_S$  on  $D_S$  compared to  $M_S$  on  $D_{ST}$  is expected. The embeddings of the source and target domains will have a portion that qualitatively overlaps and a portion that does not overlap. Our method teaches the target domain model  $M_T$  to learn the common embedding between the source and target domain. However, target domain accuracy is lost due to  $M_T$  not having access to the part of the relevant embedding that is private to the target domain.

Fig. 5(a) shows the confusion matrix from  $M_T$  on the Opportunity dataset. It can be seen that the models performed better for some activities classes like *stand* and *lie down* than for the other classes. Many *sit* samples are wrongly classified by the model as *stand*. The *other* class is commonly misclassified as *sit*, *stand*, *lie* and *walk*. Fig. 5(b) shows the corresponding ROC curves for each

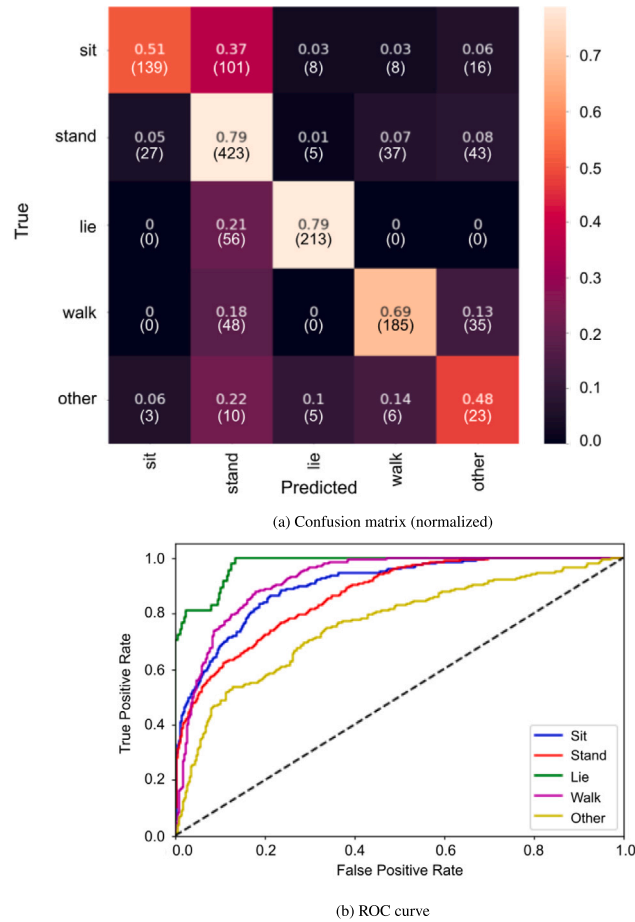


Fig. 5. (a) The confusion matrix of  $M_T$  on the Opportunity dataset highlights the difference in classification performance for different classes. The numbers in the parentheses are the absolute counts for each class and the numbers outside the parentheses are the normalized values. (b) ROC curves for each activity class in the Opportunity dataset. The classification of a class is more accurate if the overlapping information between  $e_S$  and  $e_T$  is more discriminative when compared to other classes.


of the activity classes, where the model achieved the highest AUC for *lie* while the lowest AUC for *other*. Generally,  $M_T$  would perform better on activities that are not very similar, and on activities where the common information between embeddings in the target and source domain is richer. The counterpart confusion matrices and ROC curves for the other two datasets can be found in [Appendices A](#) and [B](#), respectively.

### 6.2. Impact of quantity of training data

To get a sense of how the size of unsupervised training data  $D_{ST}$  may impact the performance of the target-domain model  $M_T$ , we trained  $M_T$  separately with 15%, 33%, 66% and 100% of the samples from  $D_{ST}$ . [Fig. 6](#) shows the F1 scores for each dataset against the number of randomly chosen samples used for this unsupervised training. The F1 score increases as the number of samples increases within each of the three datasets. As seen in the graph of the Opportunity dataset, such increase in performance seems to plateau beyond when the number of training samples exceeds a certain point. Here, the F1 performance stops improving substantially after approximately 3000 samples are included. However, as briefly discussed in [Section 6.1](#), the three datasets cannot be easily compared for performance due to key differences in the data collection process and protocol, sensors used, sensor placement, and thus the variation of difficulty in their respective activity classification.

### 6.3. Switching roles of source and target domains

We repeated the experiment [6.1](#) by switching the source and target domain body locations. That is, the torso was considered as the source body location and the wrist was considered as the target body location. This is to test our hypothesis that our model relies on the overlapping information between the source and target domains to perform classification at the target domain; hence,



**Fig. 6.** F1 scores for different number of samples from each dataset. The green stars represent MHEALTH dataset, red kdots represent PAMAP2 dataset and blue squares represent Opportunity dataset. The trend seen is that the F1 score increases as the amount of training data increases for each dataset. Comparison across datasets is not possible due to differences in factors like sensor types, sensor placements, etc.

**Table 2**

Performance metrics as in Table 1 when the source and target domains were switched, i.e. the source domain is the torso and target domain is the wrist. There was a decrease in performance of  $M_S$  on  $D_S$  which led to a decrease in performance of  $M_T$  on  $D_{ST}$  when compared to Table 1.

Metrics	Opportunity			PAMAP2			MHEALTH		
	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$
Accuracy	74.41	20.13	63.98	83.00	51.43	77.79	89.12	60.16	89.44
Precision	73.67	14.80	49.69	79.29	33.19	74.9	64.19	16.38	49.86
Recall	79.49	22.37	51.55	83.35	42.40	69.93	69.50	15.88	46.18
F1 score	76.47	17.82	<b>50.56</b>	81.27	37.24	<b>72.33</b>	66.74	16.13	<b>47.95</b>

**Table 3**

Performance metrics when M2 was trained using different loss functions for Opportunity dataset. Acc stands for accuracy, P stands for precision, R stands for recall, F1 stands for F1 score, sim stands for similarity, and reg stands for regularization. Mean absolute error loss along with L2 regularization was seen to provide the best performance.

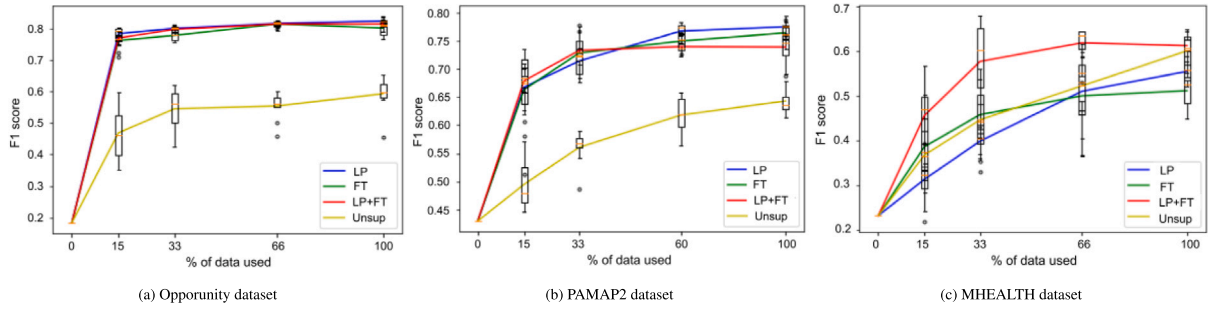
Metrics	$M_S$	Random samp	Untrained	MSE	MSLE	MAE	Cosine sim	Cosine sim + L1 reg	MAE + L1 reg	MSE + L1 reg	Cosine sim + L2 reg	MAE + L2 reg	MSE + L2 reg
Acc	76.13	26.6	24.01	52.98	54.21	49.75	59.53	60.17	64.13	62.69	57.58	<b>65.13</b>	49.10
P	73.28	20.42	14.01	44.96	47.82	46.48	54.24	57.34	59.00	66.40	54.94	<b>66.50</b>	36.99
R	79.60	20.48	26.2	50.03	50.83	51.51	56.36	57.16	64.17	59.82	57.24	<b>67.89</b>	43.27
F1	76.31	20.45	18.26	47.36	49.28	48.87	55.28	57.25	60.93	62.94	56.07	<b>67.19</b>	39.89

when the source and target domain are switched, the performance of the new classification model on the new target domain should be comparable, assuming the new  $M_S$  has similar performance as the old one. Accordingly,  $M_S$  was trained in the same supervised fashion on the torso IMU data and  $M_T$  was trained using the wrist IMU data and corresponding torso embeddings.

Table 2 shows the accuracy, precision, recall and F1 scores for this experiment 6.3 for all three datasets. The performance of  $M_S$  on  $D_S$  remained similar for the Opportunity dataset when trained on the torso. However the performance for the PAMAP2 and MHEALTH datasets dropped.  $M_S$  on  $D_S$  performs as well as or worse when trained on the torso as compared to when trained on the wrist. One possible cause is movement data loss in the torso, such as arm swinging, for certain activities. The performances of  $M_T$  follow the same performance trend as of  $M_S$ , achieving the highest performance with the PAMAP2, followed by the Opportunity and MHEALTH datasets.

#### 6.4. Testing different embedding reconstruction loss functions

To find the best loss function to be used in Eq. (6) for the activity recognition task, a variety of loss functions were tested on the Opportunity dataset for the unsupervised training of  $M_T$ . As shown in Table 3, MAE loss along with L2 regularization achieved the highest performance metrics among all the functions tested. MAE likely works well because it measures the average error between the source embeddings  $e_S$  and the transformed target embeddings  $e_T$  without overemphasizing on the point deviations in the two distributions as in some of the loss functions e.g., MSE. Furthermore, adding either L1 or L2 regularization shows performance improvement in almost all cases tested. This is expected as regularization helps reduce overfitting on the non-overlapping data between the embeddings  $e_S$  and  $e_T$ . These test results directed our choice of using MAE as the loss function for training  $M_T$  for the PAMAP2 and MHEALTH datasets, and for all the other experiments detailed in this paper.



**Fig. 7.** Comparison of F1 score obtained from proposed unsupervised method and some common transfer learning techniques as detailed in [Kumar, Raghunathan, Jones, Ma, and Liang \(2022\)](#) for different sizes of datasets when five common labels are used. LP stands for linear probing, FT for full fine tuning and LPFT for linear probing with fine tuning. It can be seen that generally transfer learning methods perform better than the proposed unsupervised method. However, this gap decreases as amount of data available decreases and the unsupervised method beats the transfer learning method at very low amounts of data.

### 6.5. Comparison with transfer learning techniques

The following experiment was conducted to answer the following two questions: (1) when is it worth labeling data in the target domain and using transfer learning techniques over our unsupervised technique of Section 3?, and (2) how much labeled data are required for transfer learning techniques to outperform our method? We tested the classification performance of three standard transfer learning techniques: linear probing (LP), full fine tuning (FT) and linear probing along with fine tuning (LPFT) in accordance with [Kumar et al. \(2022\)](#). LP involves retraining just the classification layer  $C_S$  with labeled target domain data. FT entails retraining the entire model  $M_S$  using labeled target domain data. In LPFT, the classification layer  $C_S$  is first retrained with labeled target domain data and then FT of the entire model is performed. We ran these transfer learning techniques as well as our unsupervised training method using 0%, 15%, 33%, 66% and 100% of the training data from each dataset, where 0% refers to classification without any sort of retraining, i.e. applying  $M_S$  on  $D_{ST}$ . Each of these tests was run 10 times with random sampling of the available training data in order to avoid overfitting. Performance metrics averaged over the 10 runs are reported here.

[Fig. 7](#) shows F1 scores against the percentage of data used for the three transfer learning techniques and the unsupervised domain transfer approach. It is immediately noticeable that for each approach, the performance on each of the datasets increases or plateaus as the amount of training data increases. All three transfer learning approaches outperformed our unsupervised method for all amounts of data in the Opportunity and PAMAP2 datasets; however, the performance difference between the transfer learning and the unsupervised methods was lower in the PAMAP2 dataset than in the Opportunity dataset. Note that the absolute number of samples in the PAMAP2 dataset is significantly lower than that in the Opportunity dataset. For the MHEALTH dataset, the performance of the unsupervised method is better than LP for all amounts of data, comparable to FT for 15%, 33% of the data, and better than FT for 66% and 100% of the data. Note that the MHEALTH dataset has the fewest data samples among the three datasets, suggesting that when only limited amount of data are available, there is scope for the unsupervised method to outperform some transfer learning methods. Thus, our method would outperform these transfer learning methods when very little data is available at the target body site, and these transfer learning methods would outperform our method as the amount of labeled data increases. There is a trade-off between performance and the cost of labeling data. The decision of using the unsupervised method or investing in labeling data would depend on the particular use case.

### 6.6. Using all activity classification labels in PAMAP2 and MHEALTH

Instead of just the five overlapping activity classes, we further tested our method on all the available labels in the PAMAP2 and MHEALTH. This was not done for the Opportunity as it only contains the five overlapping labels tested. As shown in [Table 4](#), the performance metrics for  $M_S$  on  $D_S$  drop for both datasets when compared to those when only 5 labels were used. This is expected because the addition of more classes makes the classification problem more difficult. However, the trend of  $M_T$  performing better than  $M_S$  on  $D_{ST}$  is still visible, suggesting that even for more challenging activity classification tasks, unsupervised training of  $M_S$  on  $D_{ST}$  helps capture overlapping information in the source and target domain embeddings. Detailed results of comparing our unsupervised method with transfer learning methods when all labels are used can be found in [Appendix D](#).

### 6.7. Comparison with STranGAN approach

STranGAN ([Faridee et al., 2022](#)) is a method to perform unsupervised domain transfer between the source and target domain using a spatial transformer. Specifically, the spatial transformer learns affine transformations on the target domain IMU in order to have discriminative features similar to that of the source domain IMU. The spatial transformer is trained together with the domain discriminator which guesses the origin of the transformed IMU signals, i.e. guesses whether the transformed IMU signals belong to the source or target domain. Thus the domain discriminator is adversarial in nature to the spatial transformer. In this experiment,

**Table 4**

Performance metrics for PAMAP2 and MHEALTH datasets when all activity labels were used. The performance of  $M_S$  on  $D_S$  decreases as compared to 1 due to increase in the number of classes. The trend of  $M_T$  on  $D_{ST}$  performing better than  $M_S$  on  $D_{ST}$  is seen in this case as well.

Metrics	PAMAP2			MHEALTH		
	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$	$M_S$ on $D_S$	$M_S$ on $D_{ST}$	$M_T$ on $D_{ST}$
Accuracy	64.66	33.55	58.50	76.99	58.35	79.05
Precision	66.57	16.54	60.96	57.39	18.68	57.03
Recall	73.66	21.87	61.78	87.09	32.12	78.16
F1 score	69.94	18.84	<b>61.37</b>	69.19	23.63	<b>65.95</b>

**Table 5**

Comparison of F1 scores for our method with the results reported for the Opportunity and PAMAP2 datasets in Faridee et al. (2022) using the STranGAN method. Our method is comparable to or performs better than (Faridee et al., 2022).

	Opportunity		PAMAP2	
	STranGAN	Ours	STranGAN	Ours
F1 score	62.6	<b>67.19</b>	<b>64.86</b>	61.37

We compared the F1 scores obtained for the Opportunity and PAMAP2 datasets with the results of STranGAN. Note that we used all the activity labels in the PAMAP2 dataset. As shown in Table 5, our method is comparable to or performs better than STranGAN for the two datasets tested. However, our method provides the key advantage of being able to use any existing activity classifier trained on the source domain, thus carrying over their useful characteristics to the target domain.

## 7. Conclusion

As activity recognition using wearable devices is gaining more applications in everyday scenarios, the flexibility of using such devices on different body locations is becoming important. In this work, we proposed an unsupervised approach to leverage an existing activity classification model, which is trained on a source body location, to perform activity classification at a target body location. We have shown that a pre-trained activity classifier can be used to perform activity classification at a new target body location if the classifier is retrained to replicate the source body location embeddings using target body location signals. The performance of the retrained classifier depends on the performance of the original classifier on the source body location. The performance of the retrained model also increases as the amount of simultaneous source-target domain data increases. Hence, the need of labeled data at the target body location is removed, significantly saving time and cost because collecting simultaneous measurements at the source and target body locations is far less expensive. Additionally, leveraging an existing activity classifier that has proven to work well on one body location allows us to retain important properties of the model, such as the model being lightweight, which is critical for running the model on edge devices like smartwatches.

## CRedit authorship contribution statement

**Chaitra Hegde:** Methodology, Software, Validation, Formal analysis, Data curation, Writing – original draft, Visualization. **Gezheng Wen:** Methodology, Software, Validation, Writing – review & editing, Visualization. **Layne C. Price:** Conceptualization, Methodology, Validation, Writing – review & editing, Supervision.

## Declaration of competing interest

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

## Data availability

The data that has been used is confidential.

## Appendix A. Confusion matrices

The confusion matrices for PAMAP2 and MHEALTH datasets are shown in Fig. 8. It can be seen that for both datasets, *other* class is classified correctly to the most extent but all the remaining classes are confused for *other* often. This might be due to the high number of *other* samples in the datasets.




Fig. 8. Confusion matrices for PAMAP2 and MHEALTH datasets.

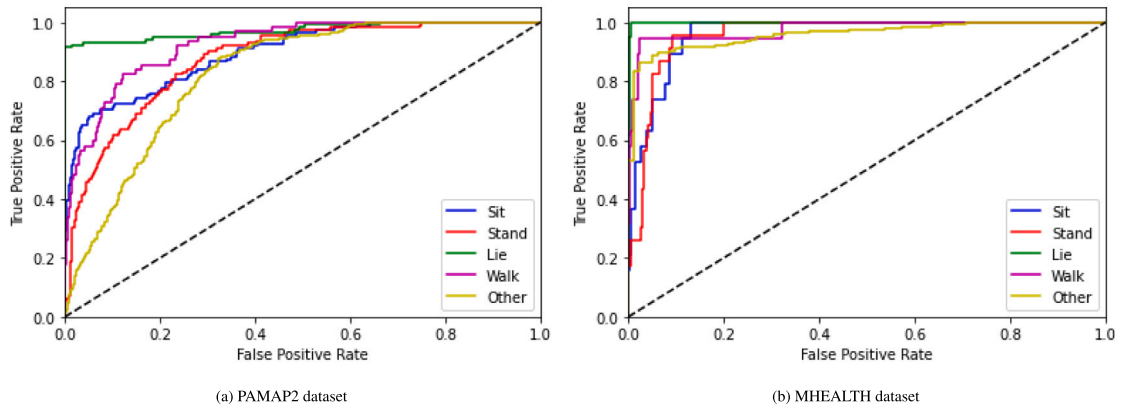


Fig. 9. ROC curves for each label for all three datasets when five labels are used.




Fig. 10. t-SNE projections for the labels *sit* and *stand* for source and target domains for all three dataset after training of  $M_T$ . The source and target projections are fairly clearly differentiated. The separation between *sit* and *stand* is a little less clear but is still visible.

### Appendix B. ROC curves

The ROC curves for PAMAP2 and MHEALTH are shown in Fig. 9. It can be seen that the ROC curves for the MHEALTH dataset, especially for the *lie* class, are near perfect. However, it is important to remember that the dataset, and test set specifically, had

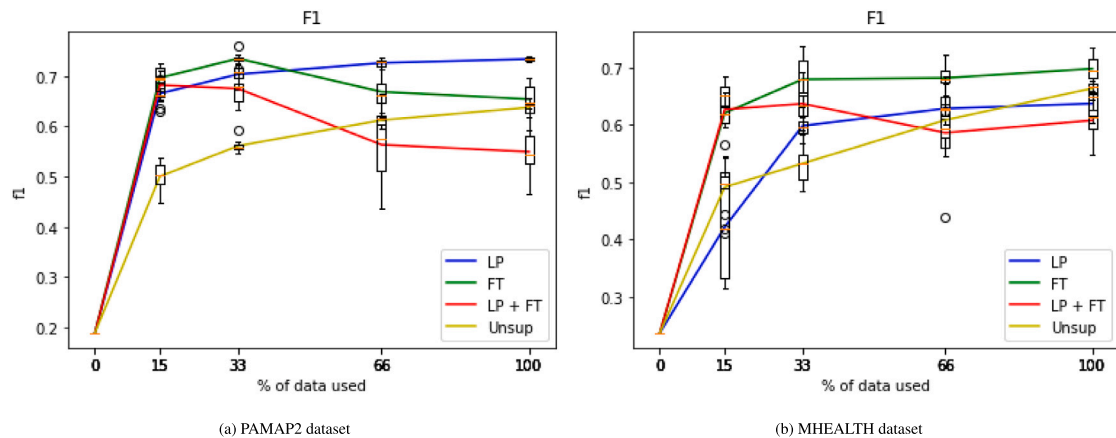


Fig. 11. Comparison of F1 score obtained from proposed unsupervised method and some common transfer learning techniques for different sizes of datasets when all labels are used for PAMAP2 and MHEALTH datasets. LP stands for linear probing, FT for full fine tuning and LPFT for linear probing with fine tuning. It can be seen that generally transfer learning methods perform better than the proposed unsupervised method. However, this gap decreases as amount of data available decreases and the unsupervised method beats the transfer learning method at very low amounts of data.

very few samples of each class other than *other*. In both datasets, *lie* class has the highest AUC and low AUC for *other* as reflected in Appendix A.

### Appendix C. t-SNE projections

The t-SNE projections for two labels, *sit* and *stand*, for source and target domains are shown in Fig. 10 for each dataset after  $M_T$  was trained. It can be seen that source and target projections are fairly clearly differentiated in all datasets. However, the separation between *sit* and *stand* is a little less clear but is still visible. It is also visible that the number of samples for MHEALTH dataset are low compared to the other two datasets.

### Appendix D. Comparison with transfer learning methods when all labels used

We compared the performance of the unsupervised method with transfer learning techniques LP, FT and LPFT as detailed in Kumar et al. (2022) when all the labels are used for PAMAP2 and MHEALTH datasets. This is shown in Fig. 11. It can be seen that generally transfer learning methods perform better than the proposed unsupervised method. However, this gap decreases as amount of data available decreases and the unsupervised method beats the transfer learning method at very low amounts of data.

## References

- Al-Nawashi, M., Al-Hazaimeh, O. M., & Saraee, M. (2017). A novel framework for intelligent surveillance system based on abnormal human activity detection in academic environments. *Neural Computing and Applications*, 28(1), 565–572.
- Alessandrini, M., Biagetti, G., Crippa, P., Falaschetti, L., & Turchetti, C. (2021). Recurrent neural network for human activity recognition in embedded systems using PPG and accelerometer data. *Electronics*, 10(14), <http://dx.doi.org/10.3390/electronics10141715>, URL: <https://www.mdpi.com/2079-9292/10/14/1715>.
- Almaadeed, N., Elharrouss, O., Al-Maadeed, S., Bouridane, A., & Beghdadi, A. (2019). A novel approach for robust multi human action recognition and summarization based on 3D convolutional neural networks. <http://dx.doi.org/10.48550/ARXIV.1907.11272>, URL: <https://arxiv.org/abs/1907.11272>.
- Banos, O., Garcia, R., Holgado-Terriza, J. A., Damas, M., Pomares, H., Rojas, I., et al. (2014). MHealthDroid: a novel framework for agile development of mobile health applications. In *International workshop on ambient assisted living* (pp. 91–98). Springer.
- Bin Morshed, M., Haresamudram, H. K., Bandaru, D., Abowd, G. D., & Ploetz, T. (2022). A personalized approach for developing a snacking detection system using earbuds in a semi-naturalistic setting. In *Proceedings of the 2022 ACM international symposium on wearable computers* (pp. 11–16). New York, NY, USA: Association for Computing Machinery, <http://dx.doi.org/10.1145/3544794.3558469>.
- Bota, P., Silva, J., Folgado, D., & Gamboa, H. (2019). A semi-automatic annotation approach for human activity recognition. *Sensors*, 19(3), <http://dx.doi.org/10.3390/s19030501>, URL: <https://www.mdpi.com/1424-8220/19/3/501>.
- Caramia, C., Torricelli, D., Schmid, M., Muñoz-Gonzalez, A., Gonzalez-Vargas, J., Grandas, F., et al. (2018). IMU-based classification of parkinson's disease from gait: A sensitivity analysis on sensor location and feature selection. *IEEE Journal of Biomedical and Health Informatics*, 22(6), 1765–1774. <http://dx.doi.org/10.1109/JBHI.2018.2865218>.
- Chen, C., Miao, Y., Lu, C. X., Xie, L., Blunsom, P., Markham, A., et al. (2019). Motiontransformer: Transferring neural inertial tracking between domains. In *Proceedings of the AAAI conference on artificial intelligence, Vol. 33* (pp. 8009–8016).
- Chen, Y., Wang, J., Huang, M., & Yu, H. (2019). Cross-position activity recognition with stratified transfer learning. *Pervasive and Mobile Computing*, 57, 1–13.
- Chen, K., Yao, L., Zhang, D., Wang, X., Chang, X., & Nie, F. (2020). A semisupervised recurrent convolutional attention model for human activity recognition. *IEEE Transactions on Neural Networks and Learning Systems*, 31(5), 1747–1756. <http://dx.doi.org/10.1109/TNNLS.2019.2927224>.
- Dobkin, B. H., Xu, X., Batalin, M., Thomas, S., & Kaiser, W. (2011). Reliability and validity of bilateral ankle accelerometer algorithms for activity recognition and walking speed after stroke. *Stroke*, 42(8), 2246–2250.

- Faridee, A. Z. M., Chakma, A., Misra, A., & Roy, N. (2022). STranGAN: Adversarially-learned spatial transformer for scalable human activity recognition. *Smart Health*, 23, Article 100226.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., et al. (2016). Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1), 2096–2030.
- Gjoreski, M., Gjoreski, H., Lustrrek, M., & Gams, M. (2016). How accurately can your wrist device recognize daily activities and detect falls? *Sensors*, 16(6), 800.
- Gjoreski, H., Lustrrek, M., & Gams, M. (2011). Accelerometer placement for posture recognition and fall detection. In *2011 seventh international conference on intelligent environments* (pp. 47–54). <http://dx.doi.org/10.1109/IE.2011.11>.
- Gleason, J., Ranjan, R., Schwarcz, S., Castillo, C. D., Cheng, J.-C., & Chellappa, R. (2018). A proposal-based solution to spatio-temporal action detection in untrimmed videos. <http://dx.doi.org/10.48550/ARXIV.1811.08496>, URL: <https://arxiv.org/abs/1811.08496>.
- Gowsikhaa, D., Abirami, S., et al. (2012). Suspicious human activity detection from surveillance videos. *International Journal on Internet & Distributed Computing Systems*, 2(2).
- Guan, Y., & Plötz, T. (2017). *Ensembles of deep LSTM learners for activity recognition using wearables*, Vol. 1. New York, NY, USA: Association for Computing Machinery, <http://dx.doi.org/10.1145/3090076>.
- Guan, D., Yuan, W., Lee, Y.-K., Gavrilov, A., & Lee, S. (2007). Activity recognition based on semi-supervised learning. In *13th IEEE international conference on embedded and real-time computing systems and applications (RTCSA 2007)* (pp. 469–475). <http://dx.doi.org/10.1109/RTCSA.2007.17>.
- Hassan, M. M., Uddin, M. Z., Mohamed, A., & Almogren, A. (2018). A robust human activity recognition system using smartphone sensors and deep learning. *Future Generation Computer Systems*, 81, 307–313. <http://dx.doi.org/10.1016/j.future.2017.11.029>, URL: <https://www.sciencedirect.com/science/article/pii/S0167739X17317351>.
- Khan, Z. N., & Ahmad, J. (2021). Attention induced multi-head convolutional neural network for human activity recognition. *Applied Soft Computing*, 110, Article 107671. <http://dx.doi.org/10.1016/j.asoc.2021.107671>, URL: <https://www.sciencedirect.com/science/article/pii/S1568494621005925>.
- Kim, M., Cho, J., Lee, S., & Jung, Y. (2019). IMU sensor-based hand gesture recognition for human-machine interfaces. *Sensors*, 19(18), 3827.
- Kim, B. H., Hong, S. H., Oh, I. W., Lee, Y. W., Kee, I. H., & Lee, S. Y. (2021). Measurement of ankle joint movements using IMUs during running. *Sensors*, 21(12), <http://dx.doi.org/10.3390/s21124240>, URL: <https://www.mdpi.com/1424-8220/21/12/4240>.
- Köping, L., Shirahama, K., & Grzegorzec, M. (2018). A general framework for sensor-based human activity recognition. *Computers in Biology and Medicine*, 95, 248–260. <http://dx.doi.org/10.1016/j.compbiomed.2017.12.025>, URL: <https://www.sciencedirect.com/science/article/pii/S0010482517304225>.
- Kumar, A., Raghunathan, A., Jones, R., Ma, T., & Liang, P. (2022). Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International conference on learning representations*. URL: <https://openreview.net/forum?id=UYneFzXSJWh>.
- Kwon, Y., Kang, K., & Bae, C. (2014). Unsupervised learning for human activity recognition using smartphone sensors. *Expert Systems with Applications*, 41(14), 6067–6074. <http://dx.doi.org/10.1016/j.eswa.2014.04.037>, URL: <https://www.sciencedirect.com/science/article/pii/S0957417414002607>.
- Lamichhane, B., Nemati, E., Ahmed, T., Rahman, M., Kuang, J., & Gao, A. (2022). A template matching based cough detection algorithm using IMU data from earbuds. In *2022 IEEE-EMBS international conference on biomedical and health informatics (BHI)* (pp. 01–04). <http://dx.doi.org/10.1109/BHI56158.2022.9926839>.
- Li, J., Kang, P., Tan, T., & Shull, P. B. (2022). Transfer learning improves accelerometer-based child activity recognition via subject-independent adult-domain adaptation. *IEEE Journal of Biomedical and Health Informatics*, 26(5), 2086–2095. <http://dx.doi.org/10.1109/JBHI.2021.3118717>.
- Li, H., Shrestha, A., Heidari, H., Le Kernec, J., & Fioranelli, F. (2020). Bi-LSTM network for multimodal continuous human activity recognition and fall detection. *IEEE Sensors Journal*, 20(3), 1191–1201. <http://dx.doi.org/10.1109/JSEN.2019.2946095>.
- Luwe, Y. J., Lee, C. P., & Lim, K. M. (2022). Wearable sensor-based human activity recognition with hybrid deep learning model. *Informatics*, 9(3), <http://dx.doi.org/10.3390/informatics9030056>, URL: <https://www.mdpi.com/2227-9709/9/3/56>.
- Ma, H., Li, W., Zhang, X., Gao, S., & Lu, S. (2019). AttnSense: Multi-level attention mechanism for multimodal human activity recognition. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19* (pp. 3109–3115). International Joint Conferences on Artificial Intelligence Organization, <http://dx.doi.org/10.24963/ijcai.2019/431>.
- Mekruksavanich, S., Jantawong, P., Hnoohom, N., & Jitpattanakul, A. (2022). Classification of physical exercise activity from ECG, PPG and IMU sensors using deep residual network. In *2022 research, invention, and innovation congress: innovative electricals and electronics (RI2C)* (pp. 130–134). <http://dx.doi.org/10.1109/RI2C56397.2022.9910287>.
- Morales, F. J. O. n., & Roggen, D. (2016). Deep convolutional feature transfer across mobile activity recognition domains, sensor modalities and locations. In *Proceedings of the 2016 ACM international symposium on wearable computers* (pp. 92–99). New York, NY, USA: Association for Computing Machinery, <http://dx.doi.org/10.1145/2971763.2971764>.
- Moya Rueda, F., Grzeszick, R., Fink, G. A., Feldhorst, S., & Ten Hompel, M. (2018). Convolutional neural networks for human activity recognition using body-worn sensors. *Informatics*, 5(2), <http://dx.doi.org/10.3390/informatics5020026>, URL: <https://www.mdpi.com/2227-9709/5/2/26>.
- Mu, F., Gu, X., Guo, Y., & Lo, B. (2020). Unsupervised domain adaptation for position-independent IMU based gait analysis. In *2020 IEEE SENSORS* (pp. 1–4). <http://dx.doi.org/10.1109/SENSORS47125.2020.9278863>.
- Nunez, J. C., Cabido, R., Pantrigo, J. J., Montemayor, A. S., & Velez, J. F. (2018). Convolutional neural networks and long short-term memory for skeleton-based human activity and hand gesture recognition. *Pattern Recognition*, 76, 80–94.
- Ogris, G., Stiefmeier, T., Lukowicz, P., & Troster, G. (2008). Using a complex multi-modal on-body sensor system for activity spotting. In *2008 12th IEEE international symposium on wearable computers* (pp. 55–62). <http://dx.doi.org/10.1109/ISWC.2008.4911585>.
- Ordóñez, F. J., & Roggen, D. (2016). Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1), <http://dx.doi.org/10.3390/s16010115>, URL: <https://www.mdpi.com/1424-8220/16/1/115>.
- Ordóñez, F. J., & Roggen, D. (2016). Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1), 115.
- Pei, Z., Cao, Z., Long, M., & Wang, J. (2018). AAAI'18/IAAI'18/AAAI'18, *Multi-adversarial domain adaptation*. AAAI Press.
- Qi, W., Su, H., Yang, C., Ferrigno, G., De Momi, E., & Aliverti, A. (2019). A fast and robust deep convolutional neural networks for complex human activity recognition using smartphone. *Sensors*, 19(17), <http://dx.doi.org/10.3390/s19173731>, URL: <https://www.mdpi.com/1424-8220/19/17/3731>.
- Rahmani, M. H., Berkvens, R., & Weyn, M. (2021). Chest-worn inertial sensors: A survey of applications and methods. *Sensors*, 21(8), <http://dx.doi.org/10.3390/s21082875>, URL: <https://www.mdpi.com/1424-8220/21/8/2875>.
- Reiss, A., & Stricker, D. (2012). Introducing a new benchmarked dataset for activity monitoring. In *2012 16th international symposium on wearable computers* (pp. 108–109). IEEE.
- Roggen, D., Calatroni, A., Rossi, M., Holleczeck, T., Förster, K., Tröster, G., et al. (2010). Collecting complex activity datasets in highly rich networked sensor environments. In *2010 seventh international conference on networked sensing systems (INSS)* (pp. 233–240). IEEE.
- Saini, R., Kumar, P., Roy, P. P., & Dogra, D. P. (2018). A novel framework of continuous human-activity recognition using kinect. *Neurocomputing*, 311, 99–111. <http://dx.doi.org/10.1016/j.neucom.2018.05.042>, URL: <https://www.sciencedirect.com/science/article/pii/S0925231218306179>.
- Stikic, M., Larlus, D., Ebert, S., & Schiele, B. (2011). Weakly supervised recognition of daily life activities with wearable sensors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12), 2521–2537. <http://dx.doi.org/10.1109/TPAMI.2011.36>.
- Stikic, M., Van Laerhoven, K., & Schiele, B. (2008). Exploring semi-supervised and active learning for activity recognition. In *2008 12th IEEE international symposium on wearable computers* (pp. 81–88). <http://dx.doi.org/10.1109/ISWC.2008.4911590>.

- Storm, F. A., Buckley, C. J., & Mazzà, C. (2016). Gait event detection in laboratory and real life settings: Accuracy of ankle and waist sensor based methods. *Gait & Posture*, 50, 42–46.
- Sukor, A. S. A., Zakaria, A., & Rahim, N. A. (2018). Activity recognition using accelerometer sensor and machine learning classifiers. In *2018 IEEE 14th international colloquium on signal processing & its applications (CSPA)* (pp. 233–238). <http://dx.doi.org/10.1109/CSPA.2018.8368718>.
- Syed, A. S., Sierra-Sosa, D., Kumar, A., & Elmaghraby, A. (2022). A deep convolutional neural network-XGB for direction and severity aware fall detection and activity recognition. *Sensors*, 22(7), 2547.
- Taha, A., Yayed, H. H., Khalifa, M., & El-Horbaty, E.-S. M. (2015). Human activity recognition for surveillance applications. In *Proceedings of the 7th international conference on information technology* (pp. 577–586).
- Badar ud din Tahir, S., Jalal, A., & Kim, K. (2021). Daily life log recognition based on automatic features for health care physical exercise via IMU sensors. In *2021 international bhurban conference on applied sciences and technologies (IBCAST)* (pp. 494–499). <http://dx.doi.org/10.1109/IBCAST51254.2021.9393204>.
- Wang, J., Chen, Y., Hu, L., Peng, X., & Yu, P. S. (2018). Stratified transfer learning for cross-domain activity recognition. In *2018 IEEE international conference on pervasive computing and communications (PerCom)* (pp. 1–10). <http://dx.doi.org/10.1109/PERCOM.2018.8444572>.
- Wang, J., Zheng, V. W., Chen, Y., & Huang, M. (2018). Deep transfer learning for cross-domain activity recognition. In *Proceedings of the 3rd international conference on crowd science and engineering*. New York, NY, USA: Association for Computing Machinery, <http://dx.doi.org/10.1145/3265689.3265705>.
- Wren, C. R., & Tapia, E. M. (2006). Toward scalable activity recognition for sensor networks. In M. Hazas, J. Krumm, & T. Strang (Eds.), *Location- and context-awareness* (pp. 168–185). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Wyatt, D., Philipose, M., & Choudhury, T. (2005). Unsupervised activity recognition using automatically mined common sense. In *Proceedings of the 20th national conference on artificial intelligence - Volume 1* (pp. 21–27). AAAI Press.
- Zhao, H., Wang, Z., Qiu, S., Shen, Y., & Wang, J. (2017). IMU-based gait analysis for rehabilitation assessment of patients with gait disorders. In *2017 4th international conference on systems and informatics (ICSAI)* (pp. 622–626). <http://dx.doi.org/10.1109/ICSAI.2017.8248364>.
- Zhao, H., Zhang, S., Wu, G., Moura, J. M., Costeira, J. P., & Gordon, G. J. (2018). Adversarial multiple source domain adaptation. *Advances in Neural Information Processing Systems*, 31.
- Zhou, X., Liang, W., Wang, K. I.-K., Wang, H., Yang, L. T., & Jin, Q. (2020). Deep-learning-enhanced human activity recognition for Internet of Healthcare Things. *IEEE Internet of Things Journal*, 7(7), 6429–6438. <http://dx.doi.org/10.1109/JIOT.2020.2985082>.
- Zhu, Q., Chen, Z., & Soh, Y. C. (2019). A novel semisupervised deep learning method for human activity recognition. *IEEE Transactions on Industrial Informatics*, 15(7), 3821–3830. <http://dx.doi.org/10.1109/TII.2018.2889315>.