

1

Contrastive Left-Right Wearable Sensors (IMUs) Consistency Matching for HAR

Dominique Nshimiyimana¹
DFKI and RPTU Kaiserslautern

Paul Lukowicz³
DFKI and RPTU Kaiserslautern

Vitor Fortes Rey²
DFKI and RPTU Kaiserslautern

Abstract

Machine learning algorithms are improving rapidly, but annotating training data remains a bottleneck for many applications. In this paper, we show how real data can be used for self-supervised learning without any transformations by taking advantage of the symmetry present in the activities. Our approach involves contrastive matching of two different sensors (left and right wrist or leg-worn IMUs) to make representations of co-occurring sensor data more similar and those of non-co-occurring sensor data more different. We test our approach on the Opportunity and MM-Fit datasets. In MM-Fit we show significant improvement over the baseline supervised and self-supervised method SimCLR, while for Opportunity there is significant improvement over the supervised baseline and slight improvement when compared to SimCLR. Moreover, our method improves supervised baselines even when using only a small amount of the data for training. Future work should explore under which conditions our method is beneficial for human activity recognition systems and other related applications.

¹nshimiyim@rptu.de

²fortes@dfki.uni-kl.de

³Paul.Lukowicz@dfki.de

1 Introduction

Machine learning algorithms are improving rapidly, but annotating training data remains a bottleneck for many applications. To overcome this problem, a current area of research focuses on learning effective sensor representations without human-annotated labels for human activity recognition [11, 27, 30, 10, 12, 8, 21, 9, 31, 22]. These methods follow the pretrain-then-finetune learning schema. In general, the encoder is pretrained in a self-supervised (SSL) way, and then the whole classifier is finetuned with a few labels. In this work, in the pretraining step (Step 1 in Figure 1), the network learns from SSL, namely the contrastive prediction task. The contrastive prediction task allows us to train the model without annotations. The resulting model is able to capture relevant characteristics of the input data and is beneficial for the downstream task [11], as mentioned in Step 2 of Figure 1. In the second step in Figure 1, the downstream task, the model needs a small amount of annotated data to train the classification model.

The success of those methods is already known in other fields such as computer vision [19, 14, 17, 2], natural language processing [13], and audio [5]. Lastly, similar methods have been proposed for human activity recognition, including contrastive predictive coding [10], multitask [27], and SimCLR [31]. The last has a Siamese schema and is based on contrastive learning. Contrastive learning trains a model to match different views of instances created from the same signal by contrasting them. The positive pair is obtained by the transformation of the input window, while the negative pair is formed by the remaining input windows of the same minibatch.

However, specific datasets may not benefit from some SSL methods. For example, in MM-Fit[29], applying the reconstruction (autoencoder) as a pretext task did not lead to better classification results. David et al. [29] proposed to learn from different five sensor modalities (fuse inertial sensor data and 3D pose) to be able to deal with the high intra-class variability found in HAR datasets. Pretrain-then-transfer is a strategy utilised in SSL in order to achieve better general representations [11], but the design or choice of suitable SSL settings is not trivial [11, 24]. Therefore, we contribute to the field by proposing a new SSL method that takes advantage of symmetry in activities by using contrastive learning between multiple sensors to tackle issues like (1) limited labelled data or (2) performance on datasets with left-right arm/leg synchronous activity.

Another advantage of contrastive learning is that it can save computational resources. For example, we do not need the preprocessing of raw data as done in [15, 6]. In addition, other methods [6] utilised handcrafted features, while our contrastive learning-based approach automatically creates useful features.

This work is encouraged by the success of pretraining techniques [11, 16, 25] and the necessity of reducing the amount of annotated data for training. We




FIGURE 1: Shows the proposed training settings. When pretraining, we use an encoder and head networks that learn to maximise the agreement of the left and right sensors using unlabelled data. The contrastive loss is utilised, in which the positive samples are left and right sensor data at time t and negatives are taken from the same batch at time $t' \neq t$. In step 1, the pretraining process takes pairs of time-synchronised data and trains the encoder in a contrastive way without any label. Next step, the encoder is finetuned using either left or right labelled data for classification. Using both left and right data produces better results. For the inference, the model needs only one device input for prediction.

then explore the effectiveness of contrastive learning methods while reducing labelled data requirements[11]. Our method contrastively matches two different sensor sources (left and right wrist or leg-worn IMUs). The training process requires only one pair of accelerometers. This work further validates the modality-agnostic character of contrastive learning and its potential for generalisation [31].

2 Related Work

The work, Contrastive Language-Image Pre-Training (CLIP) [23], proposed a technique that aligns pairs of images and texts by projecting them into the joint representation space of CLIP. Moon et al. [18] used CLIP for aligning Inertial Measurement Unit (IMU) motion sensor recordings with video and text by projecting them into the joint representation space of CLIP. In addition, contrastive learning allows research to create general-purpose training tools such as BYOL [7, 28], SimCLR [2, 31, 28], and contrastive predictive coding (CPC) [10, 19, 12, 14], just to name a few.

Contrastive Predictive Coding [10, 12] results from the combination of contrastive learning, temporal architecture networks and future timestep prediction. The positive pair is made of long-term temporal network output (context) together with future feature timesteps in the same window, while the negative input of contrastive learning is the samples from other windows of the same batch.

SimCLR [31, 28, 22] learns general representations from a dataset where a model is trained to match different view instances created from the same signal by contrasting them. The positive pair is obtained by the transformation of the input window, while the negative pair is formed by the rest of the input windows of the same batch. Then, the contrastive loss pulls the positive pair together while pulling the negative pair far away from each other in the representation space.

BYOL [28, 22] relies on two neural networks called online and target networks, that interact and learn from each other. The online network is trained to predict the target network representation of an augmented view of an image. The network is trained to minimise the mean-squared error loss between the online network’s prediction and the target representations. Only the online network is optimised via the loss, and the target network is updated via a slow-moving average of the online network.

Contrastive learning enables the learning and transfer of knowledge from one device to another [16, 25]. Previous work has focused on an arbitrary device at an arbitrary position without constraints. The success of our work lies partly in adding constraints for the position of devices, “body symmetry”.

Therefore, we focus on devices worn on the left and right sides of the body, for example, the legs or arms.

ColloSSL [16] proposes learning useful features from multiple devices by using multi-view contrastive loss. Some of the key innovations in this work are the selection of positive and negative pairs and contrastive sampling techniques. The device selection algorithm is designed to increase the likelihood of selecting ‘good’ positive and negative samples. As metrics, they compute the pairwise Maximum Mean Discrepancy (MMD) among the data points. Then, the positive pair is formed by those with the minimum MMD distance, and the rest are part of the negative samples. To compute the loss function, they applied contrastive sampling, whose role is to decide which data samples should be picked from each device for contrastive training.

Closest to our work is Learning from the Best [25], a contrastive-based method that transfers useful information from a source sensor to a target one by training a joint representation. While the objective of both works is similar –using synchronised data of two sensors to obtain a better representation – there are several differences. For example, [25] uses two separate encoders, one for each sensor, while in this work both sensors share the same one. Therefore, we also do not use translator networks to bring one representation to the other. This means that in this work, contrastive learning is applied between the encoded representations of both sensors instead of between the data and its translated version. Another difference is our contrastive loss function: we use the modified NT-Xent[31] loss, while previous work utilised the infoNCE loss. Moreover, our networks are simpler, using only 1D-convolution and fully connected layers instead of transformers combined with CNNs. While we do not compare those two methods in this work, we aim to show that even a simpler version of the contrastive learning scheme can work when taking advantage of symmetries in activities. We hope that this can contribute to position-independent HAR, as our learned model is valid for both devices without compromising performance (e.g., left- or right-hand user).

3 Method

Left-Right contrastive learning was inspired by the successful contrastive learning algorithms in computer vision and HAR, such as SimCLR [2, 31], BYOL [7, 28], learning from the best[25] algorithms, etc. The proposed method learns useful representations by maximising agreement between left- and right-worn IMU data via a contrastive loss in the latent space. In contrast to [25], our method uses one encoder for both sensors and has no translator network. Our method has the advantage of a simpler model, while [25] has the ability to train with data from any pair of devices. The training procedure illustrated in Figure 1, presents the following four major components:

1. The **left and right sensor data** synchronised within time allows building the input x of training without the need for data augmentation. The stochastic data augmentations used in other contrastive methods reflect an approximation of sensor properties. Using directly raw sensor data as views allows learning the sensor properties directly instead of through the view’s approximation by transformations.
2. An **encoder network** $encoder(\cdot)$ that predicts the representation in the latent space from the input data. The training pipeline has no restriction on architecture; a neural network of choice can be used. For example, we use the encoder from [31, 27] which has three (1D) convolutional layers.
3. A **project head** that takes the representation $h = encoder(x)$ and projects it into contrastive latent space $z = head(h)$, s. Figure 1.
4. A **contrastive loss** function that defines agreements within positive and negative samples in the same minibatch.

The proposed method can be trained using different tools of siamese architecture and contrastive loss with simple modifications such as BYOL [7], SimSiam [3], or SimCLR [2]. We chose to modify the contrastive learning framework SimCLR so that we use left and right sensor data instead of a double transformation. As in [2, 31, 28], we randomly sample a minibatch of N examples and define the contrastive loss with pairs of left x^{left} and right x^{right} sensor data derived from the minibatch. Further details can be found in Algorithm 1.

Algorithm 1 Left-Right contrastive learning.

Input: minibatch of size N , temperature constant τ

Output: \mathcal{L} loss

```

for sampled minibatch  $\{x_k\}_{k=1}^N$  do
  for all  $k \in \{1 \dots N\}$  do
     $z_{2k-1} = head(encoder(x_k^{left}))$ 
     $z_{2k} = head(encoder(x_k^{right}))$ 
  end for
  for all  $i \in \{1 \dots 2N\}$  and  $j \in \{1 \dots 2N\}$  do
     $sim_{i,j} \leftarrow$  pairwise similarity between  $z_i$  and  $z_j$ 
  end for
   $l(i, j) \leftarrow -\log \frac{\exp(sim_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k \neq i]} \exp(sim_{i,j}/\tau)}$   $\triangleright$  defines loss for a positive pair ( $i$ 
and  $j$ )
   $\mathcal{L} = \frac{1}{2N} \sum_{k=1}^N [l(2k-1, 2k) + l(2k, 2k-1)]$ 
end for

```

Note that the output of this network is normalised to lie on the unit of the hypersphere, which allows using the inner product as similarity metrics for

contrastive learning [2, 19]. The projection head is only used while training and does not participate in inference settings. In addition, the pipeline applies cosine similarity $sim(\cdot)$ as a similarity function. The remaining details are in experiments section.

4 Experiment

In this section, we present the datasets used in our experiments. First, we give a short description of our training pipeline and evaluation metrics. Finally, the results are presented.

4.1 Evaluation

We start with an overview of the two datasets used, MM-Fit and Opportunity. Afterwards, the details of the training procedure are presented along with the network’s architecture. In addition, the validation procedure is explained.




FIGURE 2: Shows MM-Fit class distribution based on windows of two seconds with 50% overlap. The preprocessed dataset has 29175 samples.

MM-Fit Dataset. An open-source dataset, MM-Fit [29], is used in our experiments. The dataset is collected from different body-worn inertial sensors and cameras (not worn), which are time-synchronised. There are four types of devices, including the *Orbbec Astra Pro* camera, the *eSense (Nokia Bell-Labs)*

earbud, the two *Mobvoi TicWatch Pro* smartwatches, and two smartphones: the *Samsung S7* and *Huawei P20*. The subject wore two smartwatches, one on each wrist, an earbud in the left ear, and smartphones in the left and right trouser pockets. The smartwatches recorded the acceleration and angular velocity at 100 Hz. In this work, we use the accelerometers from the two smartwatches. We apply windows of two seconds with a step of one second. The MM-Fit contains ten activities and 'no activity' classes; among them, we cite: squats, lunges (with dumbbells), bicep curls (alternating arms), sit-ups, push-ups, sitting overhead dumbbell triceps extensions, standing dumbbell rows, jumping jacks, sitting dumbbell shoulder presses, and dumbbell lateral shoulder raises. In addition, annotation includes repetition counting across multiple devices. For this dataset, we take into account all the described classes. Figure 2 provides a short overview of class distribution based on a window of two seconds with 50% overlap. We notice that the data are not well balanced, with the null class dominating the dataset.

Opportunity Dataset. The Opportunity dataset [26] recorded sensor data for 6 hours from four subjects in a daily living (ADL) scenario. The sensors were either integrated into the environment, or in objects, or on the body. The activities in the experiments are grouped as locomotion or gestures. This work focuses on locomotion activities where there are four classes, namely stand, walk, sit, and lie, as in [1, 25]. This dataset has more than one label type, for example, mode locomotion and mode gestures [1]. We evaluate the locomotion classes. In terms of the sensor settings, we choose acceleration data from the left and right lower arms (RLA and LLA), each having a sampling rate of 30 Hz. The null class was excluded from our experiments as, unlike in the case of MM-Fit, here we have many static classes and there is too much ambiguity between them and the null class as we are using only a single accelerometer for inference. This is also the reason we have not evaluated our method based on gestures. In other words, we follow this evaluation procedure because we focus on HAR from a single sensor, and it can be challenging to predict all classes using only a single acceleration data. For example, some activities (e.g., open, close, etc.) from the mode gesture are exclusively performed by one dominant hand [4]. Handling those classes is future work and would have to include more than one sensor as source data, as with a single one for inference, the signals to the corresponding activity would be missing.

Metrics. Three common metrics are used to study the performance of the models: macro F1, weighted F1, and accuracy. We run every experiment ten times and report the mean and standard deviation. As a baseline, we trained one model using left and right accelerations and then compute the metrics separately on left and right devices. Later, we noted 'both' to mention when left and right sensor data are used together; otherwise, it was specified whether we utilise data from the left or right arm.

Training procedure. Regarding the architecture of the networks (Figure 3), for pretraining, we need an encoder and a prediction head. The encoder has three 1D-convolutional layers with kernel of 24, 16, 8, and 32, 64, 96 filters, re-




FIGURE 3: Shows the architecture of the network. The adaptive max-pooling is applied over time. A dropout rate of 10% is used. The model is simple; the encoder and classifier together comprise 146 *thousand* parameters.

spectively. Every layer is followed by a relu activation function and a dropout of 0.1%. We applied global maximum pooling to the feature map in the third layer. In this step, we used the contrastive prediction head, which consists of three dense layers with 256, 128, and S output sizes, respectively. The contrastive loss was then optimised using SGD with a cosine decay on learning rate of 0.004, a temperature of 0.05, and a batch size of 64 over 200 epochs. Note that SimCLR was pretrained using rotation as the transformation, the same as in the code published with the paper [31].

For classification, we used the encoder with previously trained weights, followed by a classifier module. We chose a simple classifier with two fully connected layers with 1024 and class numbers respectively. The model is finetuned based on cross-entropy loss with the Adam optimiser, with a learning rate of 0.0001 for 50 epochs.

All our results refer to an unseen test set. As it is done in [29], the MM-Fit dataset is split according to the participants data: train (1, 2, 3, 4, 6, 7, 8, 16, 17, 18), validation (14, 15, 19), test set (9, 10, 11), and unseen test set (0, 5, 12, 13, 20). For Opportunity, data are preprocessed similar to [20], where we train the model on the data of all ADL, and drill sessions for the first subject and on ADL1, ADL2, ADL3 and drill sessions for Subjects 2 and 3. We report classification performance on a testing set composed of ADL4 and ADL5 for Subjects 2 and 3.

In order to evaluate the quality of learned representations, we used reduced labelled data for finetuning, namely 1, 5, 10, 50, and 100 samples per class.

The training follows the protocol in Figure 1, which implements the schema of pretrain-then-finetune. In the first step, we train the encoder in our proposed self-supervised learning manner. We trained with different hyperparameters as done in [12], then we chose the best encoder based on validation.

In the second step, the classification task, we used early stopping with patience of 5 and kept the best classification model. As in [27, 31], all layers

Method	Training	Inference	F1 Weighted	F1 Macro	Accuracy
Baseline	Left	Left	90.26±0.67	74.77±1.64	91.00±0.70
	Right	Right	90.97±0.68	76.52±2.26	91.72±0.57
	Both	Right	89.82±0.54	72.28±1.42	90.48±0.62
SimCLR[31]	Both	Left	89.93±0.54	69.58±1.76	90.69±0.43
	Both	Right	89.91±0.75	69.90±2.49	90.69±0.58
MM-Fit[29]		Left			90.74
		Right			92.72
Ours	SSL+Both	Left	93.19±0.20	82.99±0.77	93.57±0.16
		Right	93.93±0.53	85.28±1.86	94.31±0.42

TABLE 1.1: Classification (macro f1 and accuracy) results for the proposed method with comparison to baseline, MM-Fit (result reported from paper [29]) and SimCLR.

of the base encoder are frozen except for the last one, and a two-layer MLP (Multi-Layer Perceptron) classifier is added and finetuned.

4.2 Results

In Table 1.1, we can see the performance of our method when compared to the baseline (supervised training) and another SSL method (SimCLR). There we can see that for the MM-Fit dataset, our method presented around 8% macro f1 improvement compared to the baseline and even bigger gains when compared to SimCLR, which did not beat the baseline. This shows the potential of our approach, which, unlike SimCLR, does not rely on data transformations. Since MM-Fit includes symmetric motions, it is likely our method’s improvement comes directly from the lack of data transformations. One could think that having access to double the data is the source of our improvement, but our method also beats a supervised approach trained using data from both sensors as examples, which points to our contrastive training itself being beneficial.

Since MM-Fit is an unbalanced dataset, as seen in Figure 2, it is relevant to look at the confusion matrix in Figure 5 to show our performance in the most and least frequent classes. For example, in addition, the bicep curls exercise presented the lowest performance with 87.8% accuracy, which can be expected as it is the least symmetric activity. The highest performance was observed in the squats exercise, with 98.8% accuracy. The non-activity class achieved an accuracy score of 94.5%, while the remaining activities had a mean accuracy of 93.7%. This shows that the proposed method is able to handle imbalanced data.

The results on the Opportunity dataset can be read from Table 1.2, where

Method	Training	Inference	F1 Weighted	F1 Macro	Accuracy
Baseline	Left	Left	49.30±3.97	38.51±3.35	60.78±1.84
	Right	Right	43.40±3.52	34.39±6.74	54.60±3.29
	Both	Left	48.36±1.03	37.73±0.86	60.32±1.34
		Right	37.35±4.42	27.94±4.06	48.95±3.82
SimCLR[31]	Both	Left	76.17±0.85	72.56±2.00	77.02±0.80
		Right	71.08±0.81	60.57±1.15	72.24±1.30
Ours	SSL+Both	Left	77.53±0.93	73.84±1.63	78.30±0.89
		Right	72.99±0.93	62.74±1.68	74.27±1.21

TABLE 1.2: Classification results for the proposed method with comparison to baseline and SimCLR on Opportunity dataset [1, 26].

results from baseline, SimCLR, and the proposed method are compared. All SSL methods show promising improvements at all metrics of more than 25%.

In this case, our method still outperforms SimCLR, but now by less than 2% of the macro F1 score. More specifically, pretraining enhanced the results from 0.39 to 0.73 and 0.74 macro f1 for SimCLR and the proposed method, respectively (left accelerometer).

We hypothesise that our gains are smaller when compared to SimCLR due to the locomotion activities being less symmetric. Figure 8 shows that while lying and walking activities are well recognised in the Opportunity dataset, standing and sitting are the most challenging activities for the model.

4.3 Effect of a reduced training set

The Figure 4 shows results for an experiment in which we vary the number of labelled data available in the MM-Fit dataset. Here, we can see that using only a small amount of labelled data (100 samples per class), we can achieve a macro f1 score of 0.7 (dashed lines) with more than 20% improvement compared to training from scratch with the same amount of examples.

4.4 Contribution of batch size to contrastive loss

We also did an analysis of the contribution of batch size to contrastive loss. We could not observe the positive relationship between batch size and performance, in contrast to [2], where a large minibatch performs better. Figure 9 shows the contribution of batch size and latent space size to performance. As the minibatch size increases, the performance does not necessarily become better. The same behaviour is observed while comparing latent space size and accuracy.




FIGURE 4: Indicates performance of pretrained model while finetuning with reduced training set in MM-Fit (proposed settings). Supervised refers to the training from scratch (our baseline).




FIGURE 5: Shows confusion matrix on MM-Fit for unseen subjects with acceleration data from right smartwatch.

4.5 Discussion and Limitations.

The proposed method outperforms the baseline and SimCLR in all tasks. It achieves 84.03% accuracy with an improvement of 35% compared with training from scratch while having only 100 labelled samples (7,76% of the training set on MM-Fit). In addition, we observed improvements on MM-Fit and




FIGURE 6: Shows confusion matrix (left arm) on MM-Fit dataset for baseline (left) and SimCLR (right) methods.




FIGURE 7: Shows the confusion matrix on Opportunity dataset with SimCLR. The model is tested with data from left hand (left figure) and right hand (right figure).




FIGURE 8: Shows the confusion matrix on Opportunity dataset with proposed method. Model is tested with data from left hand (left figure) and right hand (right figure).

Opportunity if the entire labelled data (for trains) were available. However, the proposed method does not perform significantly better than other SSL on datasets, e.g., Opportunity, whose activities do not synchronously involve both arms or no arms, e.g., standing or sitting. Details can be visualised by

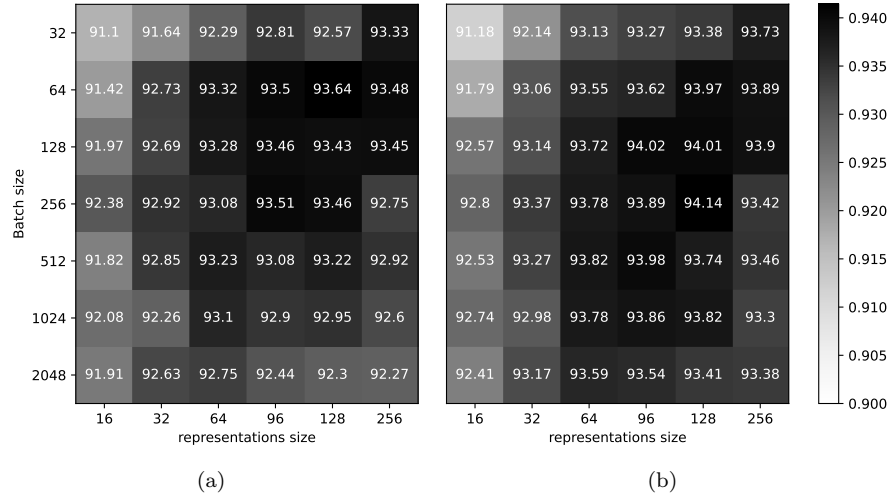


FIGURE 9: Displays the effect (accuracy in percentage) of batch size and latent space size. Figure (a) and Figure (b) report on left and right worn smartwatch, respectively.

comparing Figure 8 and Figure 7. The reason behind this may lie in the fact that the motion signals captured from the arms are perturbed by (1) independent arm movement or (2) signal propagation through different joints [25]. But there are some open questions, for example, about the performance when the dataset contains a large number of locomotion activities. Moreover, it is not always possible (or easy) to collect left-right synchronous data for various reasons, e.g., running, walking, etc.

Looking at the improvement done by the proposed method, we could observe that some classes improved only marginally from baseline. For example, the lowest improvement is observed in tricep extensions activity, where the proposed method has a 0.6% gain in macro F1 from baseline; for details, see Figures 5 and 6. In contrast to lower improvements, SimCLR classifies worse than baseline; for details, see Figure 6. Note that triceps extension in MM-Fit is symmetrical for the left and right arms. Perhaps the proposed method makes use of this synchronous behaviour, but in this case, the improvement is not significant.

5 Conclusion

The proposed pretraining technique applies contrastive learning to two different sensor data from left and right wrist-worn devices and hence does not need any transformation. In order to evaluate the performance of the learned representation for HAR, we followed the schema of 'pretrain-then-finetune'. In the results, we observed that left-right contrastive learning draws promising results. In addition, the finetuned model performed better than fully supervised learning from scratch. The results also showed the role of data representation, where the model achieved around a 0.7 macro f1 score on MM-Fit using only 100 annotated samples per class.

We have observed that our method does not need a large minibatch, in contrast to other contrastive learning frameworks based on transformations. Moreover, it can alleviate data requirements by still showing improvements given a reduced training set.

On the other hand, our proposed method was designed with symmetrical activities in mind. Thus, it requires synchronised data from both sides and may not provide substantial improvements if only one of the limbs is involved in the activity. Therefore, in future work, we will evaluate our methods with more datasets that are less symmetrical in motion.

In many areas, prior knowledge about the activities being performed is known beforehand. Thus, we will investigate how our method can be applied to those activities that fit this proposed schema. We will also investigate how to generalise the method so that other relationships between activities can be included in the contrastive framework.

Moreover, future work will also focus on studying the effects of different modalities and sensor placement, as well as data transformations.

6 Acknowledgments

The research reported in this paper was supported by the Carl Zeiss Stiftung, Germany under the Sustainable Embedded AI project (P2021-02-009).



Bibliography

- [1] Chavarriaga, R., Sagha, H., Calatroni, A., Digumarti, S.T., Tröster, G., del R. Millán, J., Roggen, D.: The opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters* **34**(15), 2033–2042 (2013). DOI <https://doi.org/10.1016/j.patrec.2012.12.014>. URL <https://www.sciencedirect.com/science/article/pii/S0167865512004205>. Smart Approaches for Human Action Recognition, Dhaka, Bangladesh
- [2] Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning, pp. 1597–1607. PMLR, Vienna, Austria (2020)
- [3] Chen, X., He, K.: Exploring simple siamese representation learning. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR, Virtual) pp. 15745–15753 (2020)
- [4] Ciliberto, M., Fortes Rey, V., Calatroni, A., Lukowicz, P., Roggen, D.: Opportunity ++: A multimodal dataset for video- and wearable, object and ambient sensors-based human activity recognition, lausanne, switzerland (2021). DOI 10.21227/vd6r-db31. URL <https://dx.doi.org/10.21227/vd6r-db31>
- [5] Georgiev, P., Bhattacharya, S., Lane, N.D., Mascolo, C.: Low-resource multi-task audio sensing for mobile and embedded devices via shared deep neural network representations. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, Hawaii, USA **1**(3) (2017). DOI 10.1145/3131895. URL <https://doi.org/10.1145/3131895>
- [6] Gjoreski, H., Kiprijanovska, I., Stankoski, S., Kalabakov, S., Broulidakis, J., Nduka, C., Gjoreski, M.: Head-AR: Human Activity Recognition with Head-Mounted IMU Using Weighted Ensemble Learning, pp. 153–167. Springer Singapore, Singapore (2021). DOI 10.1007/978-981-15-8944-7_10. URL https://doi.org/10.1007/978-981-15-8944-7_10
- [7] Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., Piot, B., Kavukcuoglu, K., Munos, R., Valko, M.: Bootstrap your own latent: A new approach to self-supervised learning. *NeurIPS*, Virtual (2020)

- [8] Haresamudram, H., Anderson, D.V., Plötz, T.: On the role of features in human activity recognition. In: Proceedings of the 2019 ACM International Symposium on Wearable Computers, ISWC '19, p. 78–88. Association for Computing Machinery, New York, NY, USA (2019). DOI 10.1145/3341163.3347727. URL <https://doi.org/10.1145/3341163.3347727>
- [9] Haresamudram, H., Beedu, A., Agrawal, V., Grady, P.L., Essa, I., Hoffman, J., Plötz, T.: Masked reconstruction based self-supervision for human activity recognition. In: Proceedings of the 2020 ACM International Symposium on Wearable Computers, ISWC '20, p. 45–49. Association for Computing Machinery, New York, NY, USA (2020). DOI 10.1145/3410531.3414306. URL <https://doi.org/10.1145/3410531.3414306>
- [10] Haresamudram, H., Essa, I., Plötz, T.: Contrastive predictive coding for human activity recognition. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., Vitual **5**(2) (2021). DOI 10.1145/3463506. URL <https://doi.org/10.1145/3463506>
- [11] Haresamudram, H., Essa, I., Plötz, T.: Assessing the state of self-supervised human activity recognition using wearables. Proceedings of the ACM (New York, NY, United States) on Interactive, Mobile, Wearable and Ubiquitous Technologies **6**(3), 1–47 (2022)
- [12] Haresamudram, H., Essa, I., Plötz, T.: Investigating enhancements to contrastive predictive coding for human activity recognition. In: 2023 IEEE International Conference on Pervasive Computing and Communications (PerCom, Atlanta , USA), pp. 232–241 (2023). DOI 10.1109/PERCOM56429.2023.10099197
- [13] Hashimoto, K., Xiong, C., Tsuruoka, Y., Socher, R.: A joint many-task model: Growing a neural network for multiple NLP tasks. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 1923–1933. Association for Computational Linguistics, Copenhagen, Denmark (2017). DOI 10.18653/v1/D17-1206. URL <https://aclanthology.org/D17-1206>
- [14] Hénaff, O.J., Srinivas, A., Fauw, J.D., Razavi, A., Doersch, C., Eslami, S.M.A., van den Oord, A.: Data-efficient image recognition with contrastive predictive coding. In: International conference on machine learning, pp. 4182–4192. PMLR, Vienna, Austria (2020)
- [15] Islam, M.S., Hossain, T., Ahad, M.A.R., Inoue, S.: Exploring Human Activities Using eSense Earable Device, pp. 169–185. Springer Singapore, Singapore (2021). DOI 10.1007/978-981-15-8944-7_11. URL https://doi.org/10.1007/978-981-15-8944-7_11

- [16] Jain, Y., Tang, C.I., Min, C., Kawsar, F., Mathur, A.: Collossl: Collaborative self-supervised learning for human activity recognition. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, Atlanta, USA **6**(1), 1–28 (2022)
- [17] Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Hawaii, USA pp. 7482–7491 (2017)
- [18] Moon, S., Madotto, A., Lin, Z., Dirafzoon, A., Saraf, A., Bearman, A., Damavandi, B.: Imu2clip: Multimodal contrastive learning for imu motion sensors from egocentric videos and text. New York, United States (2022)
- [19] Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748, New York, United States (2018)
- [20] Ordóñez, F.J., Roggen, D.: Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. Sensors, Basel, Switzerland **16**(1) (2016). DOI 10.3390/s16010115. URL <https://www.mdpi.com/1424-8220/16/1/115>
- [21] Qian, H., Pan, S.J., Miao, C.: Latent independent excitation for generalizable sensor-based cross-person activity recognition. Proceedings of the AAAI Conference on Artificial Intelligence, Virtual **35**(13), 11921–11929 (2021). DOI 10.1609/aaai.v35i13.17416. URL <https://ojs.aaai.org/index.php/AAAI/article/view/17416>
- [22] Qian, H., Tian, T., Miao, C.: What makes good contrastive learning on small-scale wearable-based tasks? Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington DC USA (2022)
- [23] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: M. Meila, T. Zhang (eds.) Proceedings of the 38th International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 139, pp. 8748–8763. PMLR, Virtual (2021). URL <https://proceedings.mlr.press/v139/radford21a.html>
- [24] Ranzato, M., Graves, A.: Unsupervised deep learning tutorial - part 2. https://ranzato.github.io/publications/tutorial_deep_unsup_learning_part2_NeurIPS2018.pdf (NeurIPS, 3 December 2018 (accessed May 1, 2023))

- [25] Rey, V.F., Suh, S., Lukowicz, P.: Learning from the best: Contrastive representations learning across sensor locations for wearable activity recognition. Proceedings of the 2022 ACM International Symposium on Wearable Computers, Hangzhou, China (2022)
- [26] Roggen, D., Calatroni, A., Rossi, M., Holleczeck, T., Förster, K., Tröster, G., Lukowicz, P., Bannach, D., Pirkl, G., Ferscha, A., Doppler, J., Holzmann, C., Kurz, M., Holl, G., Chavarriaga, R., Sagha, H., Bayati, H., Creatura, M., Millàn, J.d.R.: Collecting complex activity datasets in highly rich networked sensor environments, kassel, germany. In: 2010 Seventh International Conference on Networked Sensing Systems (INSS), pp. 233–240 (2010). DOI 10.1109/INSS.2010.5573462
- [27] Saeed, A., Ozcelebi, T., Lukkien, J.: Multi-task self-supervised learning for human activity detection. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, London, UK **3**(2), 1–30 (2019). DOI 10.1145/3328932. URL <https://doi.org/10.1145%2F3328932>
- [28] Shah, K., Spathis, D., Tang, C.I., Mascolo, C.: Evaluating contrastive learning on wearable timeseries for downstream clinical outcomes. HL4H, Virtual (2021)
- [29] Strömbäck, D., Huang, S., Radu, V.: Mm-fit: Multimodal deep learning for automatic exercise logging across sensing devices. Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., Virtual **4**(4) (2020). DOI 10.1145/3432701. URL <https://doi.org/10.1145/3432701>
- [30] Tang, C.I., Perez-Pozuelo, I., Spathis, D., Brage, S., Wareham, N., Mascolo, C.: Selfhar: Improving human activity recognition through self-training with unlabeled data. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, Virtual **5**(1), 1–30 (2021). DOI 10.1145/3448112. URL <https://doi.org/10.1145%2F3448112>
- [31] Tang, C.I., Perez-Pozuelo, I., Spathis, D., Mascolo, C.: Exploring contrastive learning in human activity recognition for healthcare. Machine Learning for Mobile Health Workshop at NeurIPS, Vancouver, Canada (2020)