

Weak-Annotation of HAR Datasets using Vision Foundation Models

Marius Bock
Ubiquitous Computing &
Computer Vision
University of Siegen
Siegen, Germany
marius.bock@uni-siegen.de

Kristof Van Laerhoven
Ubiquitous Computing
University of Siegen
Siegen, Germany
kvl@eti.uni-siegen.de

Michael Moeller
Computer Vision
University of Siegen
Siegen, Germany
michael.moeller@uni-siegen.de

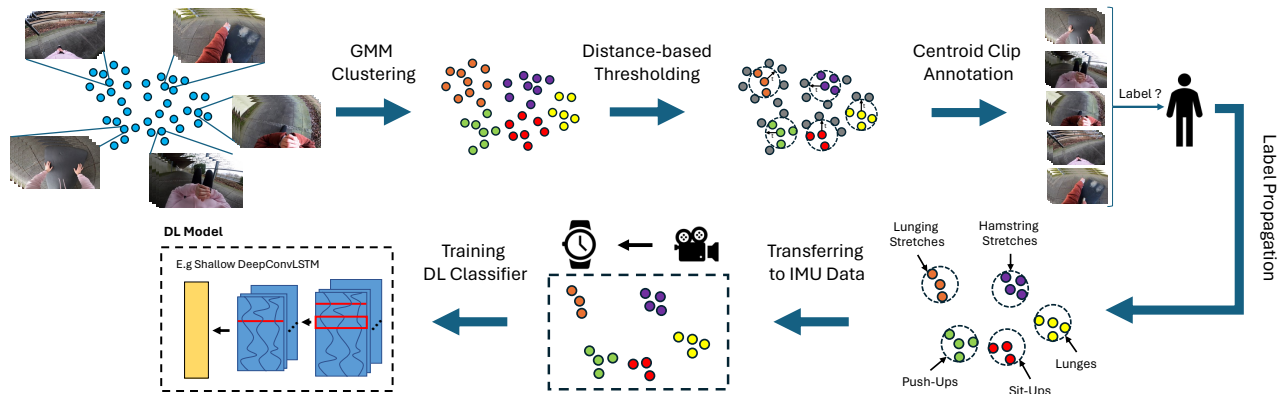


Figure 1: Our proposed weak-annotation pipeline: Visual embeddings extracted using Vision Foundation Models are clustered using Gaussian Mixture Models (GMMs). Decreasing the required labelling effort, a human annotator is only asked to annotate each cluster’s centroid video clip. Centroid labels are then propagated within each cluster. Transferred to the corresponding IMU-data, resulting weakly-annotated datasets can be used to train subsequent classifiers.

ABSTRACT

As wearable-based data annotation remains, to date, a tedious, time-consuming task requiring researchers to dedicate substantial time, benchmark datasets within the field of Human Activity Recognition in lack richness and size compared to datasets available within related fields. Recently, vision foundation models such as CLIP have gained significant attention, helping the vision community advance in finding robust, generalizable feature representations. With the majority of researchers within the wearable community relying on vision modalities to overcome the limited expressiveness of wearable data and accurately label their to-be-released benchmark datasets offline, we propose a novel, clustering-based annotation pipeline to significantly reduce the amount of data that needs to be annotated by a human annotator. We show that using our approach, the annotation of centroid clips suffices to achieve average labelling accuracies close to 90% across three publicly available HAR benchmark datasets. Using the weakly annotated datasets, we further demonstrate that we can match the accuracy scores of fully-supervised deep learning classifiers across all three benchmark datasets. Code as well as supplementary figures and results are publicly downloadable via github.com/mariusbock/weak_har.

CCS CONCEPTS

• **Human-centered computing** → Ubiquitous and mobile computing design and evaluation methods.

KEYWORDS

Data Annotation; Human Activity Recognition; Body-worn Sensors

1 INTRODUCTION

Though the automatic recognition of activities through wearable data has been identified as valuable information for numerous research fields [9], currently available wearable activity recognition benchmark datasets lack richness and size compared to datasets available within related fields. Compared with, for example, the newly released Ego4D dataset [17], it becomes apparent that currently used datasets within the inertial-based Human Activity Recognition (HAR) community are significantly smaller in terms of the number of participants, length of recordings, and variety of performed activities. One of the main drivers for this is that, even though body-worn sensing approaches allow recording large amounts of data with only minimal impact on users in various situations in daily life, wearable-based data annotation remains, to date, a tedious, time-consuming task and requires researchers to dedicate substantial time to it during data collection (taking up to 14 to 20 times longer than the actual recorded data [30]).

Following the success in other vision-related fields, researchers within the video activity recognition community have made use of feature extraction methods which provide latent representations of video clips rather than using raw image data [23, 33]. Such feature extraction methods are usually pretrained on existing large

benchmark corpora, though often not particularly related, they are capable of transferring knowledge to the activity recognition task at hand. Recently, vision foundation models [28, 29] have gained a lot of attention. Typically trained on a large amount of curated and uncurated benchmark datasets, these models have helped the community further advance in finding robust, generalizable visual feature representations.

With the majority of researchers within the wearable activity recognition community relying on the vision modality to overcome the lack expressiveness of wearable data and accurately label their to-be-released benchmark datasets offline (see e.g. [11, 18, 30]), we propose a novel annotation pipeline which makes use of visual embeddings extracted using pretrained foundation models to significantly limit the amount of data which needs to be annotated by a human annotator. Our contributions are three-fold:

- (1) We find that visual embeddings extracted using publicly-available vision foundation models can be clustered activity-wise.
- (2) We show that annotating only one clip per cluster suffices to achieve average labelling accuracies above 60% and close to 90% across three publicly available HAR benchmark datasets.
- (3) We demonstrate that using the weakly annotated datasets, one is capable of matching accuracy scores of fully-supervised deep learning classifiers across all three benchmark datasets.

2 RELATED WORK

Vision Foundation Models. The term foundation models was coined by Devlin et al. [14] and refers to models which are pretrained on a large selection of datasets. The idea of pre-training models on large benchmark datasets has been prominent within the vision community for a long time. Within the video classification community researchers demonstrated that pretrained methods such as I3D [10], VideoSwin [24] or SlowFast [15] extract discriminative feature embeddings which can be used to train subsequent classifiers. Following their success in Natural Language Processing [8, 14], researchers applied masked autoencoders on visual data input. Unlike previous methods, masked autoencoders are capable of pretraining themselves in a self-supervised manner, allowing the use of larger data sources. Two of such methods are CLIP [29] and DINOv2 [28]. The former, published by OpenAI, is a vision-language model which tries to learn the alignment between text and images. According to the authors, CLIP is pretrained on a large copora of image-text pairs scraped from the world wide web. Similarly, the recently released DINOv2 by META AI makes an effort of providing a foundation model which is capable of extracting general purpose visual features, for which the authors collected data from curated and uncurated sources.

Weakly-Supervised Wearable HAR. With the activity labelling of body-worn sensor data being a tedious task, many researchers have looked at weakly-supervised learning techniques to reduce the required amount of annotations to train subsequent classifiers. Early works such that of Stikic et al. [34] have shown to reduce the labelling efforts for training classical machine learning models through knowledge-driven approaches using graph-based label

propagation [35], multi-instance learning [36] or probabilistic methods [16, 38]. Adami and Thomaz [2] followed the works of [34], proposing an active learning framework which focuses on asking users to label only data which will gain the most classification performance boost. With the rise in popularity of Deep Learning, deep clustering algorithms have been proposed to cluster latent representations in unsupervised and semi-supervised fashion using e.g. autoencoders [4, 5, 22, 26, 42], recurrent networks [1, 16], self-supervised [31] and contrastive learning [3, 12, 20, 40, 43]. Recently, Xia et al. [43] and Tong et al. [40] demonstrated how vision-foundation models such as CLIP [29] and I3D [10] can be used to create visual, complementary embeddings to inertial data such that a contrastive loss can be calculated. This work marks one of the few instances of researchers trying to use visual data to limit the amount of annotations required in wearable activity recognition. Our work ties into the works of Tong et al. [40] and Xia et al. [43], yet we propose instead to apply vision foundation models to perform automatic label propagation between similar embeddings.

3 METHODOLOGY

3.1 Annotation Pipeline

Latent Space Clustering via Vision Foundation Models. Within the first phase we divide the unlabeled dataset into (overlapping) video clips. Given an input video stream X of a sample participant, we apply a sliding window approach which shifts over X , dividing the input data into video clips, e.g. of four second duration with a 75% overlap between consecutive windows. This process results in $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ being discretized into $t = \{0, 1, \dots, T\}$ time steps, where T is the number of windows, i.e. video clips, for each participant. Inspired by classification approaches originating from the temporal action localization community, we make use of pretrained vision foundation models to extract latent representations of each clip. That is, $\mathbf{x}_t \in \mathbb{R}^E$ represents a one-dimensional feature embedding vector associated with the video clip at time step t , where E is number of latent features the embedding vector consists of. In total we evaluated three popular pretrained foundation models: a two-stream inflated 3D-ConvNet (I3D) [10] pretrained on the RGB and optical flow features extracted from Kinetics-400 dataset [19] as well as two transformer foundation models CLIP [29] and DINOv2 [28], which were pretrained on a multitude of curated and uncurated data sources. Note that, unlike Carreira and Zisserman in [10], we use RAFT [39] instead of TV-L1 [37] optical flow estimation. As the CLIP and DINOv2 model both are not explicitly trained on optical flow features, we also test complementing embeddings of the two models by concatenating them with extracted embeddings of the inflated 3D-ConvNet trained on RAFT optical flow features of the Kinetics dataset. In order to obtain latent representations we altered models such that intermediate feature representations can be extracted. Table 1 provides details activations at which layer were considered to be the embedding of which pretrained method as well as their dimension. To merge together the frame-wise features outputted of the CLIP and DINOv2 model, we apply average pooling as detailed in [25] to obtain a single latent representation per sliding video clip.

Table 1: Network layer used for extracting embeddings of the different vision foundation models [10, 28, 29]. Subsequent layers are omitted such that the network outputs latent representations at point of the embedding layer. Note that the I3D network is used for extracting both RGB and flow features and we refer to the vision-based part of the CLIP model.

Model	Embedding Layer	Dimension
I3D	last average pool layer	\mathbb{R}^{1024}
CLIP	last projection layer (vision-CLIP)	\mathbb{R}^{768}
DINO	last layer hidden state (clf. token)	\mathbb{R}^{1024}

Having extracted latent representations of each video clip within the tested benchmark datasets, we apply Gaussian Mixture Models (GMM) to cluster the embeddings on a per-participant level. Though GMMs are not originally intended to be used with high-dimensional data, they have shown to provide good results clustering visual embeddings especially in the context of action recognition [21, 41] and allow, unlike methods such as k-nearest neighbors, more flexibility regarding the shape of clusters. Training one GMM clustering algorithm per participant and applying it to said that assigns each video clip $\mathbf{x}_t \in \mathbb{R}^E$ a cluster label $x_c \in 1, \dots, C$, where C is the number of GMM components, i.e. clusters, applied.

Weak-Labeling via Centroid Clips. Once each video clip of a study participant has been assigned a cluster label x_c , the second phase of our approach consists of a human annotator only needing to annotate one sample clip per cluster. Assuming the centroid of a cluster is most representative of all clips within that cluster, we can propagate the activity label $a \in 1, \dots, A$ of said clip to all other clips eliminating the need of annotating the other clips via a human annotator, where A is the number of activities within the dataset. As GMM do not explicitly provide a definition of a centroid of a component, we calculate the centroid clip of each cluster component being the clip which has the highest density within said cluster. That is, given the covariance matrix of each mixture component $\Sigma \in \mathbb{R}^{E \times E}$, assuming each component has its own general covariance matrix, and mean vector $\mu \in \mathbb{R}^E$, we calculate the density of each point as the logarithm of the probability density function of the multivariate normal distribution defined by μ and Σ . Having identified the centroid clip within each cluster, our approach propagates the annotation provided by the human annotator to all other clips, which were also assigned to that cluster.

As our approach forces each video clip to be assigned an activity label, we augment our clustering with a subsequent distance-based thresholding in order to remove outlier clips from the automatic labelling. Assuming that the distance of another clip to the centroid resembles its likelihood of belonging to the same activity class, we omit clips from the dataset which exceed a certain distance from their respective centroid clip, with the distance being calculated as the L^2 -norm between two embedding vectors. Even though this approach decreases the amount of data which can be used to train subsequent classification algorithms, we show to increase the overall labelling accuracy by a significant margin.

3.2 Weakly-supervised Training

Assuming inertial and video data are synchronised, we further evaluate how well the resulting annotated inertial data with non-uniform label noise is suited to be used for training inertial-based deep learning classifiers. As our benchmark algorithms of choice we use two recently published state-of-the-art methods, namely the Shallow DeepConvLSTM [6] and TinyHAR architecture [45]. We use both architectures as originally introduced by the authors, specifically using the same size and number of convolutional filters, convolutional layers and size of the recurrent layers. During training we apply a sliding window of one second with an overlap of 50%, as it proved to be provide consistent classification performances across a multitude of HAR datasets [6]. We train each network for 30 epochs using the Adam optimizer (learning rate $1e^{-4}$ and weight decay $1e^{-6}$) applying a step-wise learning rate with a decay factor of 0.9 after each 10 epochs. To mitigate the introduced label noise by our proposed weak-annotation pipeline, we calculate the loss during training using the weighted partially Huberised generalised cross-entropy (PHGCE) loss [27], which extends the definition of the generalized cross-entropy loss [44] with a variant of gradient clipping. To compare the validity of our approach, we compare amongst a set of (weakly-)annotated training approaches:

- (1) *Fully-supervised:* Fully-supervised results using the original, fully-annotated benchmark datasets.
- (2) *Few-Shot-CE:* Fully-supervised training using only the annotated clips and a weighted cross-entropy loss.
- (3) *Random-CE:* training using an equal amount of random annotated clips as in (2) and a weighted cross-entropy loss.
- (4) *Weak-CE:* Weakly-supervised training using the weakly-annotated dataset and a weighted cross-entropy loss.
- (5) *Weak-PHGCE:* Weakly-supervised training using the weakly-annotated dataset and a weighted PHGCE loss.

Table 2: Average labeling accuracy and standard deviation across study participants using different types and combinations of embeddings [10, 28, 29] extracted from three benchmark datasets [7, 13, 32] applying a GMM-based clustering using 100 clusters. Overall, a combination of CLIP and optical flow embeddings proved most consistent across all datasets.

	WEAR	Wetlab	ActionSense
(1) I3D	82.62 (± 4.65)	66.08 (± 9.53)	53.47 (± 5.95)
(2) CLIP	82.47 (± 6.03)	72.70 (± 6.42)	59.85 (± 4.42)
(3) DINOv2	79.20 (± 4.04)	69.28 (± 8.12)	60.25 (± 4.04)
(4) RAFT	76.86 (± 4.79)	51.50 (± 6.96)	45.64 (± 5.19)
(1) + (4)	85.17 (± 4.48)	60.91 (± 8.36)	53.00 (± 4.66)
(2) + (4)	83.96 (± 4.99)	66.23 (± 7.86)	57.29 (± 5.64)
(3) + (4)	79.13 (± 4.30)	70.18 (± 9.79)	56.55 (± 4.51)

3.3 Datasets

WEAR. The WEAR dataset offers both inertial and egocentric video data of 18 participants doing a variety of 18 sports-related activities, including different styles of running, stretching, and




Figure 2: Box-plot diagrams showing the distribution of labelling accuracies across study participants with increasing number of clusters. The bar plot below the box-plots provides details per cluster setting about the percentage of data compared to the total size of the three benchmark datasets [7, 13, 32] an annotator would need to annotate. One can see a clear trend that with an increase in clusters, labelling accuracy increases along with deviation across study participants decreasing.

strength-based exercises. Recordings took place at changing outdoor locations. Each study participant was equipped with a head-mounted camera and four smartwatches, one worn on each limb in a fixed orientation, which captured 3D-accelerometer data.

ActionSense. Published by DelPreto et al. [13], the ActionSense dataset provides a multitude of sensors capturing data within an indoor, artificial kitchen setup. Amongst the sensors, participants wore Inertial Measurement Units (IMUs) on both wrists as well as smart glasses which captured the ego-view of each participant. During recordings, participants were tasked to perform various kitchen chores including chopping foods, setting a table and (un)loading a dishwasher. Within their original publication, the authors provide annotations of 19 activities of 10 participants. Note that the dataset download of the ActionSense dataset provides IMU and egocentric video data of only 9 instead of 10 participants.

Wetlab. Taking place in a wetlab laboratory environment, the Wetlab dataset [32] comprises of data of 22 study participants which performed two DNA extraction experiments. For purposes of this paper we used the annotated provided by the authors of the re-occurring activities base activities (such as stirring, cutting, etc.) within the experimental protocol. During recordings, each participant wore a smartwatch in a fixed orientation on the wrist of their dominant hand, which captured 3D-accelerometer data. Unlike the WEAR and ActionSense dataset, the Wetlab dataset provides video data of a static camera which was mounted above the table at which the experiment was performed, thus capturing a birds-eye perspective of the experiment’s surroundings.

4 RESULTS

To ensure that reported performance differences are not based on statistical variance, all reported experiments are repeated three times, applying a set of three predefined random seeds. This applies both for the annotation pipeline experiments as well as weakly-supervised training results. During all annotation-based experiment mentioned in Section 4.1 we apply a clip length of four seconds along with a three second overlap between clips. We assume that

this clip length is a suitable length to be interpretable for a human annotator while simultaneously avoiding mixing multiple activities into one sliding window. Furthermore, during all experiments only the label of the centroid clip is propagated to that of all other cluster instances. Ablation experiments evaluating different clip lengths and number of annotated clips per cluster to determine the label to be propagated can be found within our code repository.

4.1 Annotation Pipeline

Table 2 shows the average labelling accuracy averaged across participants obtained when applying our proposed annotation pipeline using various types of extracted visual embeddings. One can see that in case of the WEAR [7] and ActionSense dataset [13] labelling accuracy can be improved by combining both RGB and optical flow features in case of all embeddings. Overall a combination of CLIP and optical flow features proves to be most consistent across our three benchmark datasets of choice, making it thus our embedding of choice for subsequent experiments. Applying a labelling strategy of only annotating the centroid clip of each cluster, Figure 2 presents a box-plot visualization of applying different number of clusters during the clustering of the participant-wise embeddings. One can see that by only annotating 100 clips per study participant, our proposed annotation pipeline is capable of reaching labelling accuracies above 85% in case of the WEAR and close to 70% in case of the Wetlab [32] and ActionSense dataset.

Furthermore, as evident by an overall shrinking boxplot with increasing number of clusters, our approach is becoming more stable with the standard deviation across study participants decreasing in case of all three datasets. As Figure 2 shows, applying a clustering of $C = A$, i.e. as many clusters as there are activities in the dataset, results in the clustering not being capable of differentiating the normal and complex variations of activities, different running styles and null-class from all other activities. In general, we witness a trend that by applying a larger amount of clusters than activities present in the dataset, one gives the GMM clustering enough degrees of freedom to differentiate even activities which share similarities, yet slightly differ from each other. Lastly, by

Table 3: Deep Learning results of applying two inertial-based models [6, 45] on various weakly-annotated versions of three public datasets [7, 13, 32]. Training using weakly-annotated datasets outperformed both few-shot training using only the annotated data as well as an equal amount of random annotated clips. With an increase in number of clusters our weakly-supervised approach is capable of being close to matching the predictive performance of fully-supervised baselines having manually annotated only a fraction of the actual dataset. The suffix $T-6$ ($T-4$) refer to training applying a threshold of 6 (4).

	DeepConvLSTM						TinyHAR						
	$c = 19$		$c = 50$		$c = 100$		$c = 19$		$c = 50$		$c = 100$		
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	
WEAR	<i>Fully-supervised</i>	79.89	78.36	79.89	78.36	79.89	78.36	77.83	71.89	77.83	71.89	77.83	71.89
	Few-Shot-CE	37.41	24.76	59.58	46.25	65.61	53.51	37.41	26.55	59.58	46.25	65.61	53.51
	Random-CE	45.90	31.13	59.46	46.98	65.91	53.38	23.73	23.73	59.72	46.34	66.27	55.00
	Weak-CE	42.55	34.09	64.17	54.59	73.38	63.23	49.45	38.75	66.68	54.05	71.10	59.65
	Weak-PHGCE	48.62	35.45	70.34	55.27	76.15	63.43	51.46	39.23	68.53	55.40	73.37	61.68
	Weak-CE-T-6	59.70	46.63	73.06	60.60	76.28	66.13	57.22	46.71	68.19	55.56	72.03	60.31
	Weak-PHGCE-T-6	59.39	44.97	73.17	58.84	77.55	64.77	58.78	46.27	69.47	56.29	74.05	61.67
	Weak-CE-T-4	68.86	57.33	74.72	63.93	77.81	68.22	65.68	55.64	71.31	60.16	73.93	63.42
Weak-PHGCE-T-4	61.35	47.00	74.45	60.61	76.64	64.81	62.90	50.46	72.25	60.84	74.83	63.94	
Wetlab		$c = 9$		$c = 50$		$c = 100$		$c = 9$		$c = 50$		$c = 100$	
		Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
	<i>Fully-supervised</i>	45.27	38.64	45.27	38.64	45.27	38.64	38.75	28.85	38.75	28.85	38.75	28.85
	Few-Shot-CE	15.60	11.39	21.89	16.46	22.78	17.38	15.18	11.62	22.43	16.14	25.92	18.27
	Random-CE	16.33	8.95	26.48	18.62	26.50	20.05	34.23	24.38	27.74	17.79	29.70	19.37
	Weak-CE	16.97	8.53	32.57	25.78	36.51	29.72	23.90	14.70	34.23	24.38	36.30	25.76
	Weak-PHGCE	18.62	10.48	27.64	23.17	33.78	27.62	24.06	15.19	33.79	24.35	35.53	25.53
	Weak-CE-T-6	23.01	18.12	33.78	27.08	38.41	29.77	26.14	18.63	33.79	24.21	36.20	25.76
Weak-PHGCE-T-6	20.26	13.72	28.77	23.90	32.29	26.71	25.22	18.41	33.34	24.25	34.93	25.17	
Weak-CE-T-4	23.54	18.71	32.02	26.37	35.25	28.95	25.15	17.91	32.85	23.84	34.64	25.04	
Weak-PHGCE-T-4	21.82	15.81	28.33	23.91	31.20	25.64	23.82	17.29	32.39	23.72	33.66	24.70	
ActionSense		$c = 20$		$c = 50$		$c = 100$		$c = 20$		$c = 50$		$c = 100$	
		Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
	<i>Fully-supervised</i>	20.73	14.82	20.73	14.82	20.73	14.82	22.19	18.04	22.19	18.04	22.19	18.04
	Few-Shot-CE	6.47	2.58	8.14	4.92	9.96	6.08	6.19	2.32	8.50	4.00	11.15	6.72
	Random-CE	4.69	1.63	9.42	5.09	10.02	5.77	5.01	1.70	7.80	4.19	11.44	6.50
	Weak-CE	12.32	7.67	13.91	9.94	17.35	12.20	13.74	9.75	15.58	11.55	17.16	13.45
	Weak-PHGCE	11.65	6.34	12.34	7.43	11.62	7.02	14.35	9.43	13.88	9.25	14.52	10.09
	Weak-CE-T-6	14.67	8.81	15.95	10.07	15.29	10.69	13.35	9.32	15.70	11.22	17.53	13.15
Weak-PHGCE-T-6	9.91	4.92	10.57	5.44	11.78	6.79	12.17	7.52	13.72	8.16	14.53	9.01	
Weak-CE-T-4	10.40	7.15	12.31	7.10	12.08	8.80	10.80	8.02	13.65	9.24	15.03	11.16	
Weak-PHGCE-T-4	8.34	3.86	8.92	4.35	9.91	5.56	9.24	5.82	11.07	6.09	12.11	7.35	

distance thresholding clusters and excluding instances which are exceeding a certain distance from their respective centroid, helps increasing the labelling accuracy significantly across all datasets. While a threshold of 4 helps increase the labelling accuracy well above 75% and even up to 93% in case of the WEAR dataset, the thresholding omits between 50% and up to 90% of the datasets.

4.2 Weakly-Supervised Training

As a combination of CLIP and optical-flow-based features proved to be most stable across all three datasets, we chose to use said embedding as basis for our weakly-supervised training. Table 3 provides an overview across the eight evaluated training scenarios. Our proposed weakly-supervised training is not only capable of outperforming the few-shot training using only the annotated centroid

clips, but for the case of applying 100 clusters is close to matching accuracy scores of a fully-supervised training across all three benchmark datasets, for both inertial-based architectures. Compared to a normal cross-entropy loss, the PHGCE loss provides more stable results in case of higher label noise, e.g., when not applying a distance-based thresholding and/or applying a smaller number of clusters. In general, the distance-based thresholding significantly improved results across all datasets. Although thresholding significantly reduces the amount of training data, the resulting decrease in overall labelling noise, especially for approaches that applied a lower number of clusters, improved classification results. We provide a detailed overview of the influence of thresholding on labelling accuracy and dataset size within the paper’s code repository.

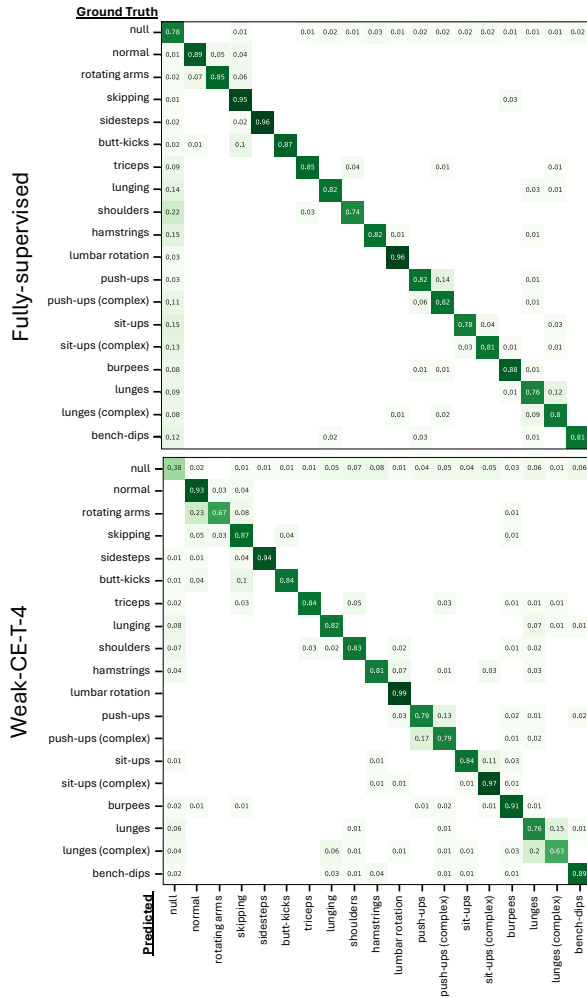


Figure 3: Confusion matrices comparing the shallow DeepConvLSTM fully-supervised results compared to that of the best performing weak-labelling approach. With exception of the NULL-class, all activities were able to be classified close to the performance of the fully-supervised approach.

Figure 3 shows a side-by-side comparison of the confusion matrices of the fully-supervised and best-performing weak-labelling approach (Weak-CE-T-4) using the WEAR dataset and shallow DeepConvLSTM. One can see that, apart from the NULL-class, classification results of all activities of the weakly-supervised training are similar to that of the fully-supervised. With the intra-class similarity of a NULL-class within datasets being quite low, it is to a larger degree grouped together with instances of other classes making it harder for an inertial-based classifier to learn unique patterns only applicable to that of the NULL-class. Looking at per-class results of the Wetlab dataset, one can see that the introduced labelling noise caused the classifier trained using weakly-annotated data is more likely to confuse activities with NULL-class rather

than vice versa. This caused the activities *transfer* and *pouring*, two classes which only have a few instances in the ground truth data and which are most frequently annotated incorrectly as NULL-class, to not be predicted correctly once across all study participants. Note that for the ActionSense dataset classification results even in the fully-supervised setting are significantly worse compared to that of the other two datasets. As evident by a nevertheless large labelling accuracy using our approach, we assume that large semantics of the dataset are too vision-centric (e.g. peeling a cucumber or a potato) to be correctly recognized using only inertial data. Nevertheless, per-class classification results of the fully-vs. weakly-supervised training show a similar confusion amongst classes, suggesting learned patterns of the classifier are similar in both training scenarios.

5 DISCUSSION & CONCLUSION

Within this paper we presented a weak-annotation pipeline for HAR datasets based on Vision Foundation Models. We showed that visual embeddings extracted using Vision Foundation Models can be clustered using Gaussian Mixture Models (GMM). Decreasing the required labelling effort, using the suggested pipeline a human annotator is only asked to annotate each cluster’s centroid video clip. By propagating the provided labels within each cluster our approach is capable of achieving average labelling accuracies above 60% and close to 90% across three popular HAR benchmark datasets. We further showed that the resulting weakly-annotated wearable datasets can be used to train subsequent deep learning classifiers with accuracy scores, in case of applying a sufficiently large number of clusters, being close to matching that of a fully-supervised training across all three benchmark datasets.

Our results underscore one of the implications recent advancements in the vision community in finding generalizable feature representations might have on the field of HAR. With the rapid progress being made in the area of foundation models, follow-ups of models such as CLIP and DINOv2 could further robustify the automatic analysis of collected video streams in wearable-based data collection. Our clustering-based pipeline thus not only improves the efficiency of data annotation but also contributes to the creation of richer and more extensive HAR benchmark datasets.

ACKNOWLEDGMENTS

We gratefully acknowledge the DFG Project WASEDO (DFG LA 275811-1) and the University of Siegen’s OMNI cluster.

REFERENCES

- [1] Alireza Abedin, Farbod Motlagh, Qinfeng Shi, Hamid RezaTofighi, and D. Ranasinghe. 2020. Towards deep clustering of human activities from wearables. <https://doi.org/10.1145/3410531.3414312>
- [2] Rebecca Adaimi and Edison Thomaz. 2019. Leveraging Active Learning and Conditional Mutual Information to Minimize Data Annotation in Human Activity Recognition. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–23. <https://doi.org/10.1145/3351228>
- [3] Abrar Ahmed, Harish Haresamudram, and Thomas Ploetz. 2022. Clustering of Human Activities from Wearables by Adopting Nearest Neighbors. In *ACM International Symposium on Wearable Computers*. <https://doi.org/10.1145/3544794.3558477>
- [4] Bandar Almaslakh, Jalal AlMuhtadi, and Abdelmonim Artoli. 2017. An Effective Deep Autoencoder Approach for Online Smartphone-Based Human Activity Recognition. *International Journal of Computer Science and Network Security* 17, 4 (2017), 160–165. http://paper.ijcsns.org/07_book/201704/20170423.pdf

- [5] Behrooz Azadi, Michael Haslgrübler, Bernhard Anzengruber-Tanase, Georgios Sopidis, and Alois Ferscha. 2024. Robust Feature Representation Using Multi-Task Learning for Human Activity Recognition. *Sensors* 24, 2 (2024), 681. <https://doi.org/10.3390/s24020681>
- [6] Marius Bock, Alexander Hoelzemann, Michael Moeller, and Kristof Van Laerhoven. 2021. Improving Deep Learning for HAR With Shallow Lstm. In *ACM International Symposium on Wearable Computers*. <https://doi.org/10.1145/3460421.3480419>
- [7] Marius Bock, Hilde Kuehne, Kristof Van Laerhoven, and Michael Moeller. 2023. WEAR: An Outdoor Sports Dataset for Wearable and Egocentric Activity Recognition. *CoRR* abs/2304.05088 (2023). <https://arxiv.org/abs/2304.05088>
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0db6fcb4967418bfb8ac142f64a-Paper.pdf
- [9] Andreas Bulling, Ulf Blanke, and Bernt Schiele. 2014. A Tutorial on Human Activity Recognition Using Body-Worn Inertial Sensors. *Comput. Surveys* 46, 3 (2014), 1–33. <https://doi.org/10.1145/2499621>
- [10] Joao Carreira and Andrew Zisserman. 2017. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/cvpr.2017.502>
- [11] Shing Chan, Hang Yuan, Catherine Tong, Aidan Acquah, Abram Schonfeldt, Jonathan Gershuny, and Aiden Doherty. 2024. CAPTURE-24: A large dataset of wrist-worn activity tracker data collected in the wild for human activity recognition. *CoRR* abs/2402.19229 (2024). <https://doi.org/10.48550/arXiv.2402.19229>
- [12] Shohreh Deldari, Hao Xue, Aaqib Saeed, Daniel V. Smith, and Flora D. Salim. 2022. COCOA: Cross Modality Contrastive Learning for Sensor Data. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–28. <https://doi.org/10.1145/3550316>
- [13] Joseph DelPreto, Chao Liu, Yiyue Luo, Michael Foshey, Yunzhu Li, Antonio Torralba, Wojciech Matusik, and Daniela Rus. 2022. ActionSense: A Multimodal Dataset and Recording Framework for Human Activities Using Wearable Sensors in a Kitchen Environment. In *Neural Information Processing Systems Track on Datasets and Benchmarks*. <https://action-sense.csail.mit.edu>
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina N. Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics*. <https://arxiv.org/abs/1810.04805>
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slow-Fast Networks for Video Recognition. In *International Conference on Computer Vision*. <https://doi.org/10.1109/ICCV.2019.00630>
- [16] Marjan Ghazvininejad, Hamid R. Rabiee, Nimma Pourdamghani, and Parisa Khanipour. 2011. HMM based semi-supervised learning for activity recognition. In *ACM International Workshop on Situation Activity & Goal Awareness*. <https://doi.org/10.1145/2030045.2030065>
- [17] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, and Xingyu Liu. 2022. Ego4D: Around the World in 3,000 Hours of Egocentric Video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/CVPR52688.2022.01842>
- [18] Alexander Hoelzemann, Julia L. Romero, Marius Bock, Kristof Van Laerhoven, and Qin Lv. 2023. Hang-Time HAR: A Benchmark Dataset for Basketball Activity Recognition Using Wrist-Worn Inertial Sensors. *MDPI Sensors* 23, 13 (2023). <https://doi.org/10.3390/s23135879>
- [19] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. 2017. The Kinetics Human Action Video Dataset. *CoRR* abs/1705.06950 (2017). <http://arxiv.org/abs/1705.06950>
- [20] Bulat Khaertdinov, Esam Ghaleb, and Stylianos Asteriadis. 2021. Contrastive Self-supervised Learning for Sensor-based Human Activity Recognition. In *IEEE International Joint Conference on Biometrics*. <https://doi.org/10.1109/IJCB52358.2021.9484410>
- [21] Anna Kukleva, Hilde Kuehne, Fadime Sener, and Jurgen Gall. 2019. Unsupervised Learning of Action Classes With Continuous Temporal Embedding. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Long Beach, CA, USA, 12058–12066. <https://doi.org/10.1109/CVPR.2019.01234>
- [22] Yongmou Li, Dianxi Shi, Bo Ding, and Dongbo Liu. 2014. Unsupervised Feature Learning for Human Activity Recognition Using Smartphone Sensors. In *Springer Second International Conference on Mining Intelligence and Knowledge Exploration*. https://doi.org/10.1007/978-3-319-13817-6_11
- [23] Xiaolong Liu, Qimeng Wang, Yao Hu, Xu Tang, Shiwei Zhang, Song Bai, and Xiang Bai. 2022. End-To-End Temporal Action Detection With Transformer. *IEEE Transactions on Image Processing* 31 (2022). <https://doi.org/10.1109/TIP.2022.3195321>
- [24] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. 2022. Video swin transformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/cvpr52688.2022.00320>
- [25] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. 2022. CLIP4Clip: An empirical study of CLIP for end to end video clip retrieval and captioning. *Neurocomputing* 508 (2022), 293–304. <https://doi.org/10.1016/j.neucom.2022.07.028>
- [26] Haojie Ma, Zhijie Zhang, Wenzhong Li, and Sanglu Lu. 2021. Unsupervised Human Activity Representation Learning with Multi-task Deep Clustering. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 1 (2021), 1–25. <https://doi.org/10.1145/3448074>
- [27] Aditya Krishna Menon, Ankit Singh Rawat, Sashank J. Reddi, and Sanjiv Kumar. 2020. Can gradient clipping mitigate label noise?. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rklB76EKPr>
- [28] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Aaleeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. 2024. DINOv2: Learning Robust Visual Features without Supervision. *CoRR* abs/2304.07193 (2024). <https://doi.org/10.48550/arXiv.2304.07193>
- [29] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. *CoRR* abs/2103.00020 (2021). <https://doi.org/10.48550/arXiv.2103.00020>
- [30] Daniel Roggen, Alberto Calatroni, Mirco Rossi, Thomas Holleczeck, Kilian Förster, Gerhard Tröster, Paul Lukowicz, David Bannach, Gerald Pirkl, Alois Ferscha, Jakob Doppler, Clemens Holzmann, Marc Kurz, Gerald Holl, Ricardo Chavarriaga, Hesam Sagh, Hamidreza Bayati, Marco Creatura, and José del R. Millán. 2010. Collecting Complex Activity Datasets in Highly Rich Networked Sensor Environments. In *IEEE Seventh International Conference on Networked Sensing Systems*. <https://doi.org/10.1109/INSS.2010.5573462>
- [31] Aaqib Saeed, Tanir Ozecebi, and Johan Lukkien. 2019. Multi-task Self-Supervised Learning for Human Activity Detection. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 2 (2019), 1–30. <https://doi.org/10.1145/3328932>
- [32] Philipp M. Scholl, Matthias Wille, and Kristof Van Laerhoven. 2015. Wearables in the Wet Lab: A Laboratory System for Capturing and Guiding Experiments. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing*. <https://doi.org/10.1145/2750858.2807547>
- [33] Dingfeng Shi, Yujie Zhong, Qiong Cao, Lin Ma, Jia Li, and Dacheng Tao. 2023. TriDet: Temporal Action Detection With Relative Boundary Modeling. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. <https://doi.org/10.1109/cvpr52729.2023.01808>
- [34] Maja Stikic, Diane Larlus, Sandra Ebert, and Bernt Schiele. 2011. Weakly Supervised Recognition of Daily Life Activities with Wearable Sensors. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 12 (2011), 2521–2537. <https://doi.org/10.1109/tpami.2011.36>
- [35] Maja Stikic, Diane Larlus, and Bernt Schiele. 2009. Multi-graph Based Semi-supervised Learning for Activity Recognition. In *IEEE International Symposium on Wearable Computers*. <https://doi.org/10.1109/ISWC.2009.24>
- [36] Maja Stikic and Bernt Schiele. 2009. Activity Recognition from Sparsely Labeled Data Using Multi-Instance Learning. In *Springer International Symposium on Location- and Context-Awareness*. https://doi.org/10.1007/978-3-642-01721-6_10
- [37] Javier Sánchez Pérez, Enric Meinhardt-Llopis, and Gabriele Facciolo. 2013. TV-L1 Optical Flow Estimation. *Image Processing On Line* 3 (2013), 137–150. <https://doi.org/10.5201/ipol.2013.26>
- [38] Jafar Tanha, Maarten Van Someren, and Hamideh Afsarmanesh. 2017. Semi-supervised self-training for decision tree classifiers. *International Journal of Machine Learning and Cybernetics* 8, 1 (2017), 355–370. <https://doi.org/10.1007/s13042-015-0328-7>
- [39] Zachary Teed and Jia Deng. 2020. RAFT: Recurrent All-Pairs Field Transforms for Optical Flow. In *European Conference on Computer Vision*. https://doi.org/10.1007/978-3-030-58536-5_24
- [40] Catherine Tong, Jinchen Ge, and Nicholas D. Lane. 2021. Zero-Shot Learning for IMU-Based Activity Recognition Using Video Embeddings. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (2021), 1–23. <https://doi.org/10.1145/3494995>
- [41] Rosaura G. VidalMata, Walter J. Scheirer, Anna Kukleva, David Cox, and Hilde Kuehne. 2021. Joint Visual-Temporal Embedding for Unsupervised Learning of Actions in Untrimmed Sequences. In *IEEE Winter Conference on Applications of Computer Vision*. <https://doi.org/10.1109/wacv48630.2021.00128>
- [42] Aiguo Wang, Guilin Chen, Cuijuan Shang, Miaofei Zhang, and Li Liu. 2016. Human Activity Recognition in a Smart Home Environment with Stacked Denoising Autoencoders. In *Web-Age Information Management*, Shaoxu Song and

- Yongxin Tong (Eds.). Vol. 9998. Springer International Publishing, Cham, 29–40. https://doi.org/10.1007/978-3-319-47121-1_3 Series Title: Lecture Notes in Computer Science.
- [43] Kang Xia, Wenzhong Li, Shiwei Gan, and Sanglu Lu. 2023. TS2ACT: Few-Shot Human Activity Sensing with Cross-Modal Co-Learning. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 4 (2023), 1–22. <https://doi.org/10.1145/3631445>
- [44] Zhilu Zhang and Mert R. Sabuncu. 2018. Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. In *Advances in Neural Information Processing Systems*. https://proceedings.neurips.cc/paper_files/paper/2018/file/f2925f97bc13ad2852a7a551802f6ea0-Paper.pdf
- [45] Yexu Zhou, Haibin Zhao, Yiran Huang, Till Riedel, Michael Hefenbrock, and Michael Beigl. 2022. TinyHAR: A Lightweight Deep Learning Model Designed for Human Activity Recognition. In *ACM International Symposium on Wearable Computers*. <https://doi.org/10.1145/3544794.3558467>