

# Evaluation of self-supervised pre-training for automatic infant movement classification using wearable movement sensors

Einari Vaaras<sup>1</sup>, Manu Airaksinen<sup>2</sup>, Sampsa Vanhatalo<sup>3</sup>, Okko Räsänen<sup>1</sup>

**Abstract**—The recently-developed infant wearable MAIJU provides a means to automatically evaluate infants’ motor performance in an objective and scalable manner in out-of-hospital settings. This information could be used for developmental research and to support clinical decision-making, such as detection of developmental problems and guiding of their therapeutic interventions. MAIJU-based analyses rely fully on the classification of infant’s posture and movement; it is hence essential to study ways to increase the accuracy of such classifications, aiming to increase the reliability and robustness of the automated analysis. Here, we investigated how self-supervised pre-training improves performance of the classifiers used for analyzing MAIJU recordings, and we studied whether performance of the classifier models is affected by context-selective quality-screening of pre-training data to exclude periods of little infant movement or with missing sensors. Our experiments show that i) pre-training the classifier with unlabeled data leads to a robust accuracy increase of subsequent classification models, and ii) selecting context-relevant pre-training data leads to substantial further improvements in the classifier performance.

**Clinical relevance**— This study showcases that self-supervised learning can be used to increase the accuracy of out-of-hospital evaluation of infants’ motor abilities via smart wearables.

## I. INTRODUCTION

Development of motor abilities in young infants (approx. 5–18 months) can be efficiently measured by the recently-developed MAIJU wearable (Motor ability Assessment of Infants with a JUmpsuit) [1, 2], which applies human activity recognition (HAR) systems via smart wearables to clinical use. MAIJU provides an objective and scalable method for automatic evaluation of infants’ motor abilities that can be used in a home environment. This information could be used for developmental research as well as to support clinical decision-making in situations like early detection of cerebral palsy or guidance of therapeutic interventions [1, 2]. As shown in [2], the accuracy of analysis produced by MAIJU is comparable to that of human evaluation from a concurrent video recording. MAIJU analysis relies on the categorical classification of infant posture, movement, and carrying, and the reliability of the final analyses depend on the accuracy of the classification tasks. Therefore, it is essential to study

ways to increase the accuracy of these classifications in order to build more robust and reliable automatic evaluation of infants’ motor abilities using movement detection, such as is present in MAIJU.

When building machine-learning models, a common scenario is that there is a scarcity of labeled data but an abundance of unlabeled data [3–8]. In self-supervised learning (SSL), a machine-learning model creates representations of input data without relying on human-annotated labels. Recently, the use of SSL for pre-training models has shown great success in various fields, such as speech processing [7–11], computer vision [6, 9–15], and natural language processing [9–11, 16].

SSL has also been successfully applied to HAR. Tang et al. [4] presented a semi-supervised method which utilizes unlabeled data via SSL to promote the performance of HAR systems. Their method achieved state-of-the-art performance with multiple HAR datasets, and the method matched the performance of similar supervised learning-based approaches with up to 10 times less labeled training data. Saeed et al. [5] proposed an SSL framework consisting of multiple auxiliary tasks for learning useful features from unlabeled raw sensory data. Their experiments demonstrated that the proposed method obtained competitive results with supervised learning-based approaches with as few as five labeled training examples for each target class.

In the present study, we investigate if SSL pre-training improves the performance of automatic infant motility assessment using the MAIJU wearable. For this, we use the classification of movement for our MAIJU data as a test case as it was the most difficult classification task in [1] and [2]. While doing so, we also seek an answer to whether classification model performance is improved by context-selective quality-screening of the data used for pre-training. The latter question is important because long-term movement recordings may include substantial amounts of data from periods that are not informative for assessing infant’s own motor performance, such as sleeping or carrying.

## II. METHODS

Fig. 1 shows an overview of the present experiments. First, a large unlabeled dataset of MAIJU data (see Sec. III-A) is used as input to SSL pre-training using three alternative data screening conditions. For SSL, we use the recently proposed data2vec [10] algorithm. After the pre-training phase, we utilize a smaller labeled dataset of MAIJU data (Sec. III-A) to fine-tune the pre-trained models. Finally, fine-tuned models are evaluated to find out if we can improve model performance

<sup>1</sup>Einari Vaaras and Okko Räsänen are with the Unit of Computing Sciences, Tampere University, Finland. Email: [einari.vaaras@tuni.fi](mailto:einari.vaaras@tuni.fi), [okko.rasanen@tuni.fi](mailto:okko.rasanen@tuni.fi)

<sup>2</sup>Manu Airaksinen is with Helsinki University Hospital, Helsinki, Finland. Email: [manu.airaksinen@hus.fi](mailto:manu.airaksinen@hus.fi)

<sup>3</sup>Sampsa Vanhatalo is with BABA Center, Pediatric Research Center, Department of Clinical Neurophysiology, New Children’s Hospital and HUS Imaging, Helsinki University Hospital, Helsinki, Finland. Email: [sampsa.vanhatalo@helsinki.fi](mailto:sampsa.vanhatalo@helsinki.fi)

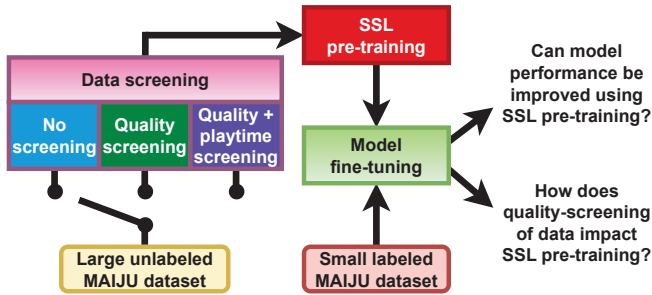


Fig. 1. An overview of the present experimental setup. Unlabeled MAIJU data is first used as input to SSL pre-training using three alternative data screening conditions, followed by fine-tuning the pre-trained models using labeled MAIJU data.

using SSL pre-training, and whether quality-screening of unlabeled data affects the benefits of SSL pre-training.

### A. Self-supervised pre-training with data2vec

Data2vec [10] is an SSL algorithm designed to work with multiple different data modalities. To achieve this, the method combines masking and predicting its own latent feature representations. As argued in [10], the power of the method comes from contextualized target representations produced by the self-attention mechanism in Transformer models. Although data2vec has been proven to work well with multiple data modalities [10, 11, 17], it should be noted that the method is not entirely data-agnostic since the method uses modality-specific feature encoders.

The data2vec method consists of two neural networks, a student network and a teacher network. Both networks have the same architecture, and they consist of a feature encoder followed by a Transformer network consisting of  $L$  Transformer encoder blocks. For the student network, the feature embeddings produced by the encoder are partially masked, whereas the teacher network receives unmasked embeddings. The student model then tries to predict the content of the masked sections of the embeddings based on the unmasked sections. The training targets are the averaged outputs of the top  $K$  Transformer encoder blocks of the teacher network, whereas the model predictions are the outputs of the last student network layer. As proposed by [10], the training targets are normalized before averaging in order to prevent representation collapse (i.e. the encoder outputs a constant time-invariant feature representation) and to prevent Transformer blocks with large output values to dominate the target features.

The weights of the teacher network are an exponentially moving average (EMA) [18] of the student network weights, i.e. after each minibatch update of the student network

$$\theta_t = \tau\theta_t + (1 - \tau)\theta_s, \quad (1)$$

where  $\theta_t$  are the weights of the teacher network,  $\theta_s$  are the weights of the student network, and  $\tau$  is a hyperparameter controlling the rate at which the weights of the teacher

network are updated. The value of  $\tau$  changes during training from a value of  $\tau_0$  into  $\tau_{end}$  linearly in  $\tau_n$  minibatch updates.

## III. EXPERIMENTS

### A. Data

We conducted our pre-training experiments using a large unlabeled dataset of MAIJU data (172 recordings, 62 distinct participants, 1333 h, referred to as “DS1” from now on), and for our fine-tuning experiments we utilized a smaller labeled dataset of MAIJU data [2] (41 recordings and distinct participants, 29.3 h, referred to as “DS2”). Most of the recordings come from typically developing infants. There are 11 recordings (3 participants) in DS1 and 3 recordings (3 participants) in DS2 with diagnosed motor disabilities. For both datasets, each sensor from all four limbs produced data in 3 accelerometer and 3 gyroscope channels (x, y, and z) with a sampling frequency of 52 Hz, resulting in 24-channel data. Following [1, 2, 19], the raw signals were pre-processed by interpolating the data packets received from each sensor (via a Bluetooth connection) into a common 52-Hz time-stamp base, removing gyroscope bias, and applying temporal smoothing with a seven-tap median filter. The pre-processed raw signals were then windowed into 120-sample frames (2.3 seconds) with an overlap of 50%.

DS1 contains multiple-hour to day-long recordings of 4–20 month-old infants obtained from their homes with a minimally controlled recording setup. In addition, the data comes with parental reports of which segments of the long recording consisted of free-play (autonomous active movement) of the baby, here referred to as “playtime”. Since DS1 was gathered in a minimally-controlled manner, there is no guarantee that the baby has been awake or let alone whether the MAIJU jumpsuit has been on the baby outside of the reported playtime. Furthermore, since the raw sensor data has been streamed over a Bluetooth connection, it is highly probable that one or multiple sensors lose connection occasionally. Consequently, without any context-selective quality-screening, a high proportion of the data is so-called “flat data” in which there is very little signal variation over time. This flat data is not informative for assessing infants’ motor performance, and therefore it is relevant to examine whether including or discarding the flat data in DS1 via quality-screening affects pre-training, as described in Sec. III-B.

DS2 contains up to one hour of data per recording, and it was obtained within a semi-controlled recording setup with parallel video at the infants’ homes ( $N = 17$  recordings) or at the lab in a home-like setting ( $N = 24$ ). The data contains nine annotated movement categories for each 2.3-s frame: still, roll left, roll right, pivot left, pivot right, proto movement, elementary movement, fluent movement, and transition. Each movement category was annotated by either two ( $N = 9$ ) or three ( $N = 32$ ) domain experts, and the data was quality-screened by the annotators during the annotation process. Table I gives further details on the number of frames and the total duration in minutes for each annotated movement category. We use all labeled data for the model training process, but for testing our trained models we only use the

TABLE I

ANNOTATED MOVEMENT CATEGORIES AND THEIR FRAME COUNTS FOR DS2 DATASET USED FOR MODEL FINE-TUNING AND TESTING. FOR EACH MOVEMENT CATEGORY, THE NUMBER OF FRAMES, THE TOTAL DURATION IN MINUTES, AND THE PERCENTAGE OF FRAMES COMPARED TO THE TOTAL NUMBER OF CATEGORIZED DATA IS GIVEN. THE NUMBER OF FRAMES WHICH ARE NOT QUALIFIED FOR MODEL FINE-TUNING AND TESTING (BABY OUT OF SCREEN OR BEING HANDLED BY CAREGIVER) IS ALSO REPORTED.

Movement category	# frames / duration (% of categorized data)
Still	38046 / 732 min (48.8%)
Roll left	627 / 12 min (0.8%)
Roll right	601 / 12 min (0.8%)
Pivot left	996 / 19 min (1.3%)
Pivot right	1103 / 21 min (1.4%)
Proto movement	20609 / 396 min (26.5%)
Elementary movement	5656 / 109 min (7.3%)
Fluent movement	5500 / 106 min (7.1%)
Transition	4753 / 91 min (6.1%)
Baby out of screen / being handled by caregiver	13558 / 261 min

frames in which the three annotators agreed on the label. For training our models, we used the so-called iterative annotation refinement (IAR) labels from [2]. The use of IAR labels was shown in [1] to yield better classification performance compared to the original annotator ratings.

### B. Model pre-training

For data2vec pre-training, the feature encoder was the same convolutional neural network (CNN) sensor encoder module that was used in [2]. The module combines the raw accelerometer and gyroscope signals at a frame-level and outputs 160-dimensional latent representations for each input frame. The only modification to the CNN encoder was the addition of layer normalization [20] after the last two convolution operations in order to make the pre-training process more stable. The Transformer model consisted of  $L = 12$  Transformer encoder blocks with an input and output dimensionality of 160, a feed-forward inner dimensionality of 640, and 10 attention heads in each multi-head self-attention. As in the original data2vec implementation [10], a Gaussian error linear unit (GeLU) [21] activation function was used in the Transformer model, and a linear projection layer consisting of 160 units was added both after the encoder and the Transformer. Instead of using absolute positional encodings with sinusoids (as in e.g. [22]), a relative positional encoder was used in the present experiments. Similar to e.g. [7, 10], this relative positional encoder was a 1D convolution (kernel size = 13, padding = 6, stride = 1) followed by a GeLU activation and layer normalization.

For data2vec pre-training, DS1 was randomly split into a training and validation set in a ratio of 80:20 sequences of 260 consecutive frames each (i.e. 5 minutes). End-of-recording

sequences shorter than 260 frames were padded into a length of 260 frames, and these padded frames were then ignored by the data2vec model during pre-training. Pre-training was tested using three data screening conditions:

- 1) *No screening*: Utilizing all unlabeled data (16,081 sequences, 1333 hours).
- 2) *Quality screening*: Screening out all the sequences in which over 20% of the frames have more than one sensor dropped out (13,119 sequences, 1087 hours).
- 3) *Quality + playtime screening*: Screening out all the sequences in which over 20% of the frames have more than one sensor dropped out and/or the data was not marked as playtime (4,669 sequences, 387 hours).

Each minibatch consisted of 32 sequences (160 minutes of signal). Mean squared error loss and Adam [23] optimizer were used. The loss was only computed for the masked frames. As training targets, we use the instance-normalized [24] and averaged outputs of all Transformer encoder blocks ( $K = 12$ ) of the teacher model. Similar to [10], we use the outputs of the feed-forward parts of each Transformer block as the outputs of our teacher model. Early stopping with a patience of 25 epochs based on the validation loss was used, and the weights of the best student model according to the validation loss were used for fine-tuning (Sec. III-C). An initial learning rate of  $10^{-4}$  was used with a reduction factor of 0.5 based on the validation loss with a patience of 15 epochs.

For masking the embeddings of the student network, we randomly select frames in each minibatch with a probability of  $p = 0.15$  to be mask starting indices. For each mask starting index we mask the embedding of that frame and the next two frames, resulting in a minimum mask length of three frames. Mask spans can overlap, which results in approximately 38% of each sequence to be masked on average. Contrary to [10], we do not replace the masked embeddings with a learnable mask token, but instead we simply replace the masked embeddings with zeros. For the EMA update of the teacher weights, we use  $\tau_0 = 0.9998$ ,  $\tau_{end} = 0.99999$ , and  $\tau_n = 10,000$ . As in [10], the weights of the feature encoder and the positional encoder were shared between the student and teacher networks, and therefore the EMA update only concerned the Transformer part of the network.

Each pre-training experiment was run both with and without data augmentation (see [19] for further details on the data augmentation methods). The augmentation methods that were used were sample dropout ( $p = 0.1$ ), sensor dropout ( $p = 0.1$ ), and sensor rotation ( $p = 0.3$ ). For sensor rotation, we utilized the Euler rotation formula where the yaw, pitch, and roll angles were sampled randomly from a uniform distribution between  $\pm 15^\circ$ . Data augmentation was only applied to the training samples. As noted in e.g. [15], the use of data augmentations may play a crucial role in SSL pre-training.

During pilot experiments, we noticed that representation collapse was very likely to occur during pre-training when a high proportion of input data was flat. These cases were when playtime screening was not utilized and when data augmentation was used with a too high probability for sample

or sensor dropout. Hence, for our pre-training experiments, sample and sensor dropout probabilities were kept at a relatively low value ( $p = 0.1$ ) when data augmentation was used, and pre-training was restarted if a representation collapse occurred.

### C. Model fine-tuning with movement category labels

DS2 was used for fine-tuning the pre-trained models (Sec. III-B) for the infant movement classification task. The models were trained using recording-level 10-fold cross-validation in which each training fold was split on a recording-level so that approximately 80% of the training fold recordings were used for training and the rest for validation. The classification model produces a prediction of the infant’s movement category for every input frame. Each training was run three times, once with a full training and validation set, and also with randomly selecting 5% or 15% of the training and validation set samples to simulate having smaller labeled datasets. This information could be utilized when building classifiers from scratch for other similar domains. At test time, an aggregate confusion matrix was calculated across the folds, followed by unweighted average F1 (UAF1) score calculation from the matrix.

The pre-trained models were fine-tuned in two stages. In the first stage, two randomly initialized fully-connected exponential linear unit (ELU) [25] layers followed by a softmax function were added after the Transformer encoder blocks in order to turn the Transformer output into categorical probabilities for each movement category. These two ELU layers consisted of 160 and 9 units, with a dropout of 40% after the first layer. These layers were fine-tuned separately by freezing the weights of the encoder and Transformer parts of the network. Weighted categorical cross-entropy loss, Adam optimizer, and a batch size of 260 consecutive frames were used. Substantial class imbalance was handled during training by weighting classes in loss calculation by their inverse frequency. Early stopping with a patience of 100 epochs based on the UAF1 score of the validation set was used to select the best-performing model. An initial learning rate of  $4 \cdot 10^{-5}$  was used with a reduction factor of 0.5 based on the UAF1 score of the validation set with a patience of 30 epochs.

In the second stage, the entire network was fine-tuned with the same training hyperparameters as in the first stage of the training process with three exceptions: First, at the beginning of the training, the learning rate was linearly increased from  $4 \cdot 10^{-8}$  to  $4 \cdot 10^{-5}$  in 20 epochs, after which the same learning rate scheduler was used as in the first stage. Second, to match the model training in [2], sample dropout ( $p = 0.3$ ) and sensor dropout ( $p = 0.3$ ) were applied randomly to the training samples as a means of data augmentation in order to make the trained models more robust. Third, a dropout of 40% was used in the Transformer part of the network, both after the self-attention blocks and also after both fully-connected layers in the feed-forward blocks.

During our experiments, we observed that by adding randomly-initialized classification layers to the end of the

TABLE II

THE RESULTS OF THE PRESENT EXPERIMENTS. THE HIGHEST UAF1 (%) SCORE FOR EACH TRAINING DATA SAMPLING SIZE IS **BOLDED**. THE RESULTS WITH AN UAF1 SCORE LOWER THAN THE WAVE NET BASELINE ARE MARKED AS **RED**, AND THE RESULTS WITH AN UAF1 SCORE HIGHER THAN THE BASELINE ARE MARKED AS **BLUE**.

		Pre-training settings		Training data sampling size		
		Data screening	Data augmentation	5%	15%	100%
Baseline		–	–	55.4	69.4	80.6
Transformer		–	–	21.7	40.1	80.6
Pre-trained models	Transformer	No screening	No	38.5	46.2	80.5
	Transformer	No screening	Yes	37.9	46.0	80.6
	Transformer	Quality	No	43.8	53.3	80.9
	Transformer	Quality	Yes	44.4	55.0	81.0
	Transformer	Quality + playtime	No	57.1	74.8	81.9
	Transformer	Quality + playtime	Yes	59.6	75.6	82.1

network and then fine-tuning the entire model at once nullifies the benefit of pre-training. We also observed that the same adverse effect occurs if the learning rate is too high at the start of the second fine-tuning stage.

### D. Classification baseline

As a baseline for our experiments, we used the same WaveNet-based [26] classification model that was used in [1, 2, 19]. WaveNet is a CNN-based model that combines causal filters with gated dilated convolutions, and it has been shown to work particularly well with MAIJU data (see e.g. [19]). Furthermore, in order to better compare the benefit of pre-training, we include a non-pretrained (randomly initialized) Transformer model baseline in our experiments. This Transformer model has the same architecture as the fine-tuned data2vec Transformer models in Sec. III-C. Both the WaveNet and the Transformer baselines were directly trained with DS2 using the same training hyperparameters as in the second pre-training stage, with the exception of not having the learning rate linearly increased at the beginning of training. Since the data2vec method requires having a Transformer-based architecture, we do not experiment with pre-training the WaveNet model. However, it should be noted that there are also SSL methods that could be used to pre-train the WaveNet model, such as [9] and [13].

## IV. RESULTS

Results of the experiments are presented in Table II. The results demonstrate that overall the baseline model, WaveNet, is very competitive performance-wise compared to the Transformer model, even though the Transformer model has nearly 34 times the number of trainable parameters of the WaveNet model. With training data sampling sizes of 5% and 15%, it is evident that the number of training samples is too small for the randomly initialized Transformer model since the model’s performance is notably below that of the WaveNet model. In these cases the WaveNet model excels, most probably due to the fact that it does not overfit to the training

True label \ Predicted label	Still	Roll L	Roll R	Pivot L	Pivot R	Proto	Elementary	Fluent	Transition
Still	0.92	0.00	0.00	0.00	0.00	0.05	0.01	0.00	0.01
Roll L	0.02	0.91	0.01	0.00	0.01	0.02	0.00	0.00	0.01
Roll R	0.01	0.01	0.90	0.00	0.00	0.07	0.00	0.00	0.01
Pivot L	0.01	0.00	0.00	0.93	0.00	0.04	0.01	0.00	0.00
Pivot R	0.02	0.00	0.00	0.00	0.90	0.06	0.02	0.00	0.00
Proto	0.10	0.01	0.00	0.01	0.01	0.78	0.05	0.00	0.04
Elementary	0.01	0.00	0.00	0.01	0.01	0.06	0.73	0.08	0.10
Fluent	0.01	0.00	0.00	0.00	0.00	0.00	0.08	0.89	0.03
Transition	0.01	0.00	0.00	0.00	0.00	0.05	0.06	0.02	0.86

Fig. 2. The confusion matrix normalized for recall values for the best-performing classification model, computed as an aggregate confusion matrix across all folds in the recording-level 10-fold cross-validation.

data as easily as the Transformer model when the number of training samples is small and when the models do not have any a priori knowledge of the data. However, utilizing pre-training without any data screening or with quality screening gives a large performance boost compared to using a non-pretrained model, albeit this increase in performance does not make the Transformer models outperform the WaveNet model which did not utilize pre-training. Only when using the most high-quality data for pre-training with “quality + playtime” screening does the Transformer model outperform the WaveNet model.

When using all labeled data for training (100% sampling size), there was no practical difference between the performance of the WaveNet model, the randomly initialized Transformer model, and the pre-trained model when no data screening was used during pre-training. Using quality screening gave a small performance boost of 0.4 percentage points in terms of UAF1 score, while using “quality + playtime” screening gave the best overall performance with a UAF1 score of 82.1%. Data augmentation during pre-training helped in achieving a better performance level, but only when data screening was used. Overall, the results indicate that pre-training is essential for models with a large number of trainable parameters, in particular if the number of labeled training samples is small. The results also show that data quality screening plays a crucial role in pre-training since a better classification performance was always achieved with it, other conditions being equal.

The aggregate confusion matrix for the best-performing classification model (pre-trained Transformer with “quality + playtime” screening and data augmentation) is shown in Fig. 2. Overall, the classifier performs well across all nine movement categories, as denoted by the high recall values on the diagonal of the matrix. The most difficult categories to

predict were “proto movement” and “elementary movement”, for which the former got often misclassified as “still” and the latter as either “fluent movement” or “transition”. The corresponding mean precision and recall, as derived from the confusion matrix, were 78.3% and 86.9%, respectively.

One potential explaining factor for the differences observed due to data screening could be attributed to the variance of the embeddings produced by the feature encoder. We observed that when less data screening was used, the proportion of flat data was higher in the pre-training dataset. This, on the other hand, caused the feature encoder to learn a feature representation with very little variance over time, and consequently, the Transformer model has an easier task in predicting masked latent feature representations in this low-variance feature space. In contrast, we observed that with more data screening, the feature encoder learned a richer feature representation in terms of variance over time, which also means that the Transformer model has to solve a more complex task when it tries to predict masked latent feature representations in a higher-variance feature space. An additional cause might simply be that with more data screening, the pre-training data represents the labeled fine-tuning data more closely.

## V. CONCLUSIONS

In the present study, we investigated whether SSL pre-training can improve the performance of automatic assessment of infants’ movement from wearable recordings with the MAIJU jumpsuit; additionally, we assessed how much context-selective quality-screening of pre-training data may improve model fine-tuning performance. Our results show that pre-training is crucial for models with a large number of trainable parameters, especially if the amount of labeled training data is small. Our results also show that screening of pre-training data is important and it leads to substantial increase in model performance; the effect is particularly clear when the pre-training data is selected to only include data with direct relevance for the eventual classifier tasks. These findings could be used to further encourage research of SSL methods for HAR and clinical healthcare applications. Potential future work includes further developing multimodal SSL algorithms, such as the data2vec algorithm used in the present experiments, to be more robust against representation collapse during SSL pre-training.

## ACKNOWLEDGMENT

O.R. was supported by Academy of Finland (AKA) grants no. 345365 and 314602. E.V. was supported by AKA grant no. 335872 and a grant from Finnish Brain Foundation. M.A. was supported by AKA grants no. 314573, 335872, and 343498. S.V. was supported by AKA grants no. 314602, 335788, 335872, 332017, and 343498, Finnish Pediatric Foundation (Lastentautiensäätiö), Finnish Brain Foundation, Sigrid Juselius Foundation, and HUS Children’s Hospital/HUS diagnostic center research funds.

## REFERENCES

- [1] M. Airaksinen, O. Räsänen, E. Ilén, T. Häyrynen, A. Kivi, V. Marchi, A. Gallen, S. Blom, A. Varhe, N. Kaartinen, L. Haataja, and S. Vanhatalo, "Automatic Posture and Movement Tracking of Infants with Wearable Movement Sensors," *Scientific Reports*, vol. 10, no. 169, 2020.
- [2] M. Airaksinen, A. Gallen, A. Kivi, P. Vijaykrishnan, T. Häyrynen, E. Ilén, O. Räsänen, L. M. Haataja, and S. Vanhatalo, "Intelligent wearable allows out-of-the-lab tracking of developing motor abilities in infants," *Communications Medicine*, vol. 2, no. 69, 2022.
- [3] H. Haresamudram, I. Essa, and T. Plötz, "Assessing the state of self-supervised human activity recognition using wearables," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 6, no. 3, p. 1–47, 2022.
- [4] C. I. Tang, I. Perez-Pozuelo, D. Spathis, S. Brage, N. Wareham, and C. Mascolo, "SelfHAR: Improving Human Activity Recognition through Self-Training with Unlabeled Data," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 5, no. 1, 2021.
- [5] A. Saeed, V. Ungureanu, and B. Gfeller, "Sense and Learn: Self-supervision for omnipresent sensors," *Machine Learning with Applications*, vol. 6, no. 100152, 2021.
- [6] H.-Y. Lee, J.-B. Huang, M. Singh, and M.-H. Yang, "Unsupervised Representation Learning by Sorting Sequences," in *Proc. ICCV*, 2017, pp. 667–676.
- [7] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations," in *Proc. NeurIPS*, 2020, pp. 12 449–12 460.
- [8] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 29, p. 3451–3460, 2021.
- [9] A. van den Oord, Y. Li, and O. Vinyals, "Representation Learning with Contrastive Predictive Coding," *arXiv preprint arXiv: 1807.03748*, 2018.
- [10] A. Baevski, W. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language," in *Proc. ICML*, 2022, pp. 1298–1312.
- [11] A. Baevski, A. Babu, W.-N. Hsu, and M. Auli, "Efficient Self-supervised Learning with Contextualized Target Representations for Vision, Speech and Language," *arXiv preprint arXiv: 2212.07525*, 2022.
- [12] O. J. Henaff, A. Srinivas, J. D. Fauw, A. Razavi, C. Doersch, S. M. A. Eslami, and A. van den Oord, "Data-Efficient Image Recognition with Contrastive Predictive Coding," in *Proc. ICLR*, 2020.
- [13] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. A. Pires, Z. D. Guo, M. G. Azar, B. Piot, K. Kavukcuoglu, R. Munos, and M. Valko, "Bootstrap Your Own Latent: a New Approach to Self-Supervised Learning," in *Proc. NeurIPS*, 2020, pp. 21 271–21 284.
- [14] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked Autoencoders Are Scalable Vision Learners," in *Proc. CVPR*, 2022, pp. 16 000–16 009.
- [15] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A Simple Framework for Contrastive Learning of Visual Representations," in *Proc. ICML*, 2020, pp. 1597–1607.
- [16] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
- [17] F. Agbavor and H. Liang, "Artificial Intelligence-Enabled End-To-End Detection and Assessment of Alzheimer's Disease Using Voice," *Brain Sciences*, vol. 13, no. 1, 2023.
- [18] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proc. NeurIPS*, 2017, pp. 1195–1204.
- [19] M. Airaksinen, S. Vanhatalo, and O. Räsänen, "Comparison of end-to-end neural network architectures and data augmentation methods for automatic infant motility assessment using wearable sensors," *Sensors*, vol. 23, no. 7, 2023.
- [20] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer Normalization," *arXiv preprint arXiv: 1607.06450*, 2016.
- [21] D. Hendrycks and K. Gimpel, "Gaussian Error Linear Units (GELUs)," *arXiv preprint arXiv: 1606.08415*, 2016.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Proc. NeurIPS*, 2017, pp. 6000–6010.
- [23] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proc. ICLR*, 2015.
- [24] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance Normalization: The Missing Ingredient for Fast Stylization," *arXiv preprint arXiv: 1607.08022*, 2016.
- [25] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)," in *Proc. ICLR*, 2016.
- [26] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "WaveNet: A Generative Model for Raw Audio," *arXiv preprint arXiv: 1609.03499*, 2016.