
Utilizing Expert Features for Contrastive Learning of Time-Series Representations

Manuel Nonnenmacher^{*12} Lukas Oldenburg¹ Ingo Steinwart² David Reeb¹

Abstract

We present an approach that incorporates expert knowledge for time-series representation learning. Our method employs expert features to replace the commonly used data transformations in previous contrastive learning approaches. We do this since time-series data frequently stems from the industrial or medical field where expert features are often available from domain experts, while transformations are generally elusive for time-series data. We start by proposing two properties that useful time-series representations should fulfill and show that current representation learning approaches do not ensure these properties. We therefore devise ExpCLR, a novel contrastive learning approach built on an objective that utilizes expert features to encourage both properties for the learned representation. Finally, we demonstrate on three real-world time-series datasets that ExpCLR surpasses several state-of-the-art methods for both unsupervised and semi-supervised representation learning.

1. Introduction

Contrastive learning (CL) has led to significant advances in the field of representation learning (Bromley et al., 1993; Hadsell et al., 2006). Recently, the introduction of the NCE-loss (Gutmann & Hyvärinen, 2010) and its variants (Oord et al., 2018; Chen et al., 2020; He et al., 2020; Grill et al., 2020; Zbontar et al., 2021) pushed the state-of-the-art for representation learning forward. These loss functions make use of different views of the data, which are transformations of the data that leave the individual label of each sample

invariant. During training, by minimizing the loss function, the different views of the same sample are then pushed towards the same representations. The aim is often to learn representations where samples with the same label are close and samples with different labels are well separated.

Such an approach is particularly successful in the vision domain, where different views (or transformations) of the data are readily available. Contrary to that, in other data domains such as time-series data, these transformations are not as easily found or requiring deep task and domain knowledge. Further, purely relying on the different views of the data does not generally lead to good representations of the data for time-series (Iwana & Uchida, 2021; Eldele et al., 2021). On the other hand, time-series datasets often stem from an industrial or medical domain, where domain experts can usually provide expert features. Expert features are therefore often readily available while access to good data transformations is generally not.

Therefore, the goal of this work is to make use of these expert features in order to find good embeddings or representations of time-series datasets utilizing expert features instead of data transformations. The use of expert features introduces several new challenges such as, which loss-function one should use and how to define good expert features. Apart from these challenges the use of expert features also circumvents several problems present in transformation based CL approaches such as the sampling bias for negative samples and the challenge of finding good transformations of the data, which can be critical for CL (Tian et al., 2020; Tamkin et al., 2020).

While expert features can usually be provided by domain experts, it remains open to find a new loss function to train using these expert features. One major challenge is that these expert features are usually not discrete class labels but continuous features, such as temperature, speed or energy values. These expert features could stem from additional measurements, (potentially expensive) simulations, or be calculated from expert mappings that use domain knowledge, directly. Therefore, using existing loss functions such as the NCE-loss or the pair-loss is not straightforwardly possible. To guide the design of a loss function, we will first establish two properties that a representations based

¹Bosch Center for Artificial Intelligence (BCAI), Robert Bosch GmbH, Renningen, Germany ²Institute for Stochastics and Applications, University of Stuttgart, Stuttgart, Germany. Correspondence to: Manuel Nonnenmacher <manuel.nonnenmacher@de.bosch.com>.

on expert features should fulfill to be useful for a range of downstream tasks. First, we want the representations of two input samples with similar expert features to be close to each other; and second, we want a pair of samples with very different expert features to be far apart in the representation space. These two principles will help us design a loss function that encourages a representation with desirable properties.

While a representation or embedding can be used for a variety of downstream tasks (i.e. predictive models, outlier detection, active learning, identifying similar samples), the two properties defined earlier assure that the representation is suited for all of these tasks. On the other hand, good performance on one of these tasks does not assure that the two properties are fulfilled (see Sec. 3.3) and therefore does also not give any performance guarantees for any of the other tasks.

Our contributions are as follows:

- We introduce two properties that a useful data representation based on expert features should satisfy, and show that only excelling in a single downstream task does not necessarily lead to a representation having the desired properties.
- Next, we introduce a novel loss function which is able to utilize continuous expert features to learn a useful representation of the data. We show that the representation obtained by minimizing our loss function attains both previously defined properties. We name the method utilizing this novel loss function ExpCLR.
- Finally, we compare our approach to several state-of-the-art approaches in the unsupervised and semi-supervised setting. For each of the three real-world datasets, our method outperforms or is on par with all other methods for several evaluation metrics.

PyTorch code implementing our method is provided at <https://github.com/boschresearch/expclr>.

2. Related Work

Our work draws on existing literature including self-supervised representation learning, knowledge distillation, metric learning and works that try to combine expert knowledge with neural networks.

While our loss is inspired by the pair-loss function (Hadsell et al., 2006; Le-Khac et al., 2020) it makes use of components used by many other contrastive loss functions (Gutmann & Hyvärinen, 2010; Chen et al., 2020; Zbontar et al., 2021; Grill et al., 2020; He et al., 2020). Although having a similar objective to ours (Wang & Isola, 2020), these

loss functions use augmentations to create positive label pairs, while we make use of expert features to determine the similarity of pairs allowing the use of continuous labels or features. Further, some of these works suffer from the negative sample bias (Chuang et al., 2020), a problem which our loss function does not encounter. Additionally, the loss function proposed by us can also be used with discrete labels in a setting similar to supervised CL (Khosla et al., 2020).

Most of the CL approaches mentioned in the previous section are applied in the vision domain, where most of the progress was made. Recently, an increasing number of works apply CL to time-series data. An early work was Contrastive Predictive Coding (CPC) (Oord et al., 2018). More recent works try to combine classical CL approaches with time-series specific training objectives and augmentations such as slicing (Tonekaboni et al., 2020; Franceschi et al., 2019; Zheng et al., 2021), forecasting (Eldele et al., 2021) and neural processes (Kallidromitis et al., 2021).

Two fields that are also closely linked to our approach are Deep Metric Learning and Knowledge Distillation, where a small student network tries to imitate the output or metric induced by a larger teacher network (Park et al., 2019; Kim et al., 2020). While traditional deep metric learning does not make use of the inherent continuous nature of the teacher model, recent works have tried go beyond binary supervision and make use of this (Kim et al., 2019; 2021). Also closely linked to this field and our work is the field of Knowledge Distillation (Gou et al., 2021), especially the works of Park et al. (2019) and Yu et al. (2019), which also take geometric relations of the teacher model into account. While these works use similar loss functions, which can also make use of the continuous nature of the expert features, our loss function leads to superior performance (Sec. 4). Further, we have a different goal, which is to precondition our representations by using the expert features to obtain a representation with favorable properties, which can then be used for a multitude of downstream tasks.

Lastly our method tries to incorporate expert knowledge with neural network training. There are several other works aiming to achieve this (Chattha et al., 2019; Hu et al., 2016). The two most relevant works are SleepPriorCL (Zhang et al., 2021) and TREBA (Sun et al., 2021). Similar to our work, TREBA and SleepPriorCL try to learn an embedding for trajectories to improve labeling-efficiency. In contrast to our work, TREBA not only proposes a contrastive loss to handle continuous expert features but further combines contrastive learning with several other training objectives such as reconstruction and consistency. While both SleepPriorCL and TREBA aim to use expert features to create pseudo-labels for unsupervised and semi-supervised representation learning by discretizing the continuous expert features, ExpCLR is able to assess expert feature distances continuously into

its objective; this avoids information loss coming from the binary positive vs. negative grouping.

3. Method

3.1. Contrastive Learning Setting

A neural network encoder (or simply *encoder*) maps samples from the input domain $x \in \mathbb{R}^{c \times T}$, where c denotes the number of input channels and T the number of time steps of each sample¹, to an embedding or *representation* $E_\theta(x) \in \mathbb{R}^e$. The encoder’s parameters (weights) are $\theta \in \mathbb{R}^P$, which are updated in supervised learning by minimizing a loss-function $\mathcal{L}(E_\theta(X), Y)$ on the training set $(X, Y) = \{(x_1, y_1), \dots, (x_N, y_N)\}$, where the labels y_i are often discrete classes $y \in \{1, \dots, C\}$.

In *contrastive representation learning* the loss function is chosen in such a way that representations of samples with the same class label are pulled closer to each other w.r.t. the Euclidean norm, while samples with different class labels are pushed away from each other. This can be achieved by minimizing a contrastive loss function, e.g. the “triplet-Loss” (Chechik et al., 2010) or the “NCE-Loss” (Gutmann & Hyvärinen, 2010). Another prominent contrastive loss is the so called *pair-loss* function (Hadsell et al., 2006; Le-Khac et al., 2020):

$$\begin{aligned} \mathcal{L}_{pair}(E(X), Y) & \quad (1) \\ &= \frac{1}{N^2} \sum_{i,j=1}^N s_{ij} D_{ij}^2 + \max \{0, (1 - s_{ij})^2 \Delta^2 - D_{ij}^2\}, \end{aligned}$$

where s_{ij} is the discrete similarity measure defined by $s_{ij} := \delta_{y_i=y_j}$, $\Delta \in \mathbb{R}_+$ is a hyperparameter, and $D_{ij} := \|E(x_i) - E(x_j)\|_2$ denotes the Euclidean distance. While the first term in the sum is responsible for pulling closer together the representations of similar points, i.e. pairs with the same labels, the second term aims to push representations of dissimilar pairs to a distance of at least Δ . So far, this is the *supervised setting* of CL, where all labels are provided. Next, we describe the *semi-supervised setting*, where labels are provided only for a fraction of the dataset, and the *unsupervised setting* as the extreme case with no labels available.

In the unsupervised setting, most CL algorithms make use of transformations $\{m_1, \dots, m_V\}$ with $m_v : \mathbb{R}^{c \times T} \rightarrow \mathbb{R}^{c \times T}$, which leave the class label invariant. These transformations can then be used to create so-called “views” $x_i^v := m_v(x_i^0)$ of a data sample $x_i^0 := x_i$. By assumption, all x_i^v have the same y -label as the original x_i^0 (even if this label is unknown) and therefore discrete similarity measure $s_{i,i}^{0,v} = 1$.

¹For notational simplicity we consider here time-series of fixed length, while our formalism is straightforward to extend to varying-length time-series.

A number of other randomly selected data samples x_j^0 ($j \neq i$), e.g. the other samples in the batch, are then considered negative samples with similarity $s_{i,j}^{0,0} = 0$, such that one can write down a CL loss like (Eq. 1) using these transformations. The expert knowledge used for unsupervised learning is thus the transformations m_v leaving class labels invariant. The training pushes the encoder towards being invariant w.r.t. the transformations m_v . In the vision domain, a great number of sensible transformations m_v are known (cropping, rotation, translation, etc.).

In the time-series domain, however, finding invariant transformations is less intuitive, and one can easily be led astray (Iwana & Uchida, 2021). But since many time-series datasets come from fields like industry or medicine in particular, expert features are often readily available from domain experts. These expert features may be discrete or (more usually) continuous, and could e.g. be calculated from the input time-series by an expert mapping $f : \mathbb{R}^{c \times T} \rightarrow \mathbb{R}^d$, i.e. $f_i := f(x_i)$, and/or be available from additional measurement sensors on the training data. In our work we assume to be given a set of expert features² $F = \{f_1, \dots, f_N\}$ for the training inputs X . The full dataset to train our embedding is thus (X, F, Y) , where Y may contain label information for any percentage of input data points, ranging between the supervised and the fully unsupervised setting.

3.2. Desiderata for Representations

We aim to employ the given expert features F in a way to learn a good representation of the input time-series. To see how to best utilize F , we first discuss what properties a useful representation should have w.r.t. the expert features. We follow the general ideal behind CL, which is to push points with the same (resp. different) labels together (resp. apart). For continuous-valued expert features $f_i \in \mathbb{R}^d$, this motivates an encoder E with the following properties:

- (P1) If expert features of two points are similar, i.e. $\|f_i - f_j\|_2$ is small, then $\|E(x_i) - E(x_j)\|_2$ should also be small, i.e. the corresponding representations be similar.
- (P2) If $\|f_i - f_j\|_2$ is large, then $\|E(x_i) - E(x_j)\|_2$ should also be large.

Here, x_i, x_j are any two samples from X , and f_i, f_j their given expert features. While both properties can be important for predictive models, (P1) is especially important for outlier detection, while (P2) is important for identifying similar samples and (safe) active learning.

The question may arise why one cannot directly use the expert features f as the representation $E(x)$, which would

²Note, we do not assume the expert mapping f to be given. It is unavailable e.g. in the HAR dataset (Sec. 4).

satisfy both properties trivially. To start, note that we do *not* assume the full expert mapping f to be available, but merely the features f_i for the given inputs X ; i.e. one couldn't evaluate $E(x_{test})$ at test inputs with such a prescription. And even if f were available, downstream tasks often benefit from fine-tuning the encoding function $E = E_\theta$ further (Sec. 4.5), which is generally not possible or successful with the expert feature mapping f . Second, the mapping f would not allow to freely choose the dimension e for the representation space, but bind it to the feature dimension d . Finally, we observe in experiments that a learned E allows to exceed the performance over the original features f in downstream tasks even for the unsupervised setting (Tab. 2).

We formalize the properties (P1) and (P2) by defining *bilipschitz representations* w.r.t. the given set of expert features:

Definition 1 (bilipschitz representation). A representation $E : \mathbb{R}^{c \times T} \rightarrow \mathbb{R}^e$ is called a $[l_-, l_+]$ -bilipschitz representation for $0 < l_- \leq l_+ < \infty$ if $\forall i, j \in \{1, \dots, N\}$:

$$l_- \|E_i - E_j\|_2 \leq \|f_i - f_j\|_2 \leq l_+ \|E_i - E_j\|_2,$$

where $E_i := E(x_i)$ and $E_j := E(x_j)$.

Note that we require (and are able to evaluate) the condition in Def. 1 only on the training set X and *not* on all potential input points $x \in \mathbb{R}^{c \times T}$. But from this, one can derive statistical bounds on the pair-Lipschitz constant $\|f - f'\|_2 / \|E(x) - E(x')\|_2$ for test points x, x' (App. D), even though it is generally impossible to satisfy Def. 1 on an infinite set of inputs when the feature dimension $d > e$ exceeds the representation dimension.

The larger l_- and the smaller l_+ is in Def. 1, the better are the guarantees one can provide for (P1) and (P2), respectively. In the ideal case we have $l_- = l_+$, i.e. the Euclidean distance in the representation space is proportional to the distance of the expert features.

3.3. Learning Representations via Feedforward Models

Having established the properties a useful representation should have, we may ask how one can obtain such a representation from the expert features F . A first approach might be to put the expert features into bins to arrive again at discrete class labels, and then proceed with a standard contrastive loss function such as the pair-loss. Such an approach cannot generally lead to a bilipschitz representation; in particular, it cannot provide guarantees on l_- and l_+ due to arbitrariness in choosing the bins and the absence of a relative distance measure between different bins. A similar problem occurs with other methods that generate pseudo-labels such as SleepPriorCL (Zhang et al., 2021).

As an alternative approach, one might add a linear layer $M = M_\phi$ on top of the encoder E_θ and jointly train θ and ϕ to predict the given expert features, e.g. by minimizing the

MSE-loss. However, such a procedure does not necessarily lead to good guarantees for the two properties (P1), (P2) we are trying to fulfill, as the following shows:

Proposition 1. Let $\mathcal{L}_{mse}(f', f) = \|f' - f\|_2^2$ be the MSE-loss, E_θ the encoder, and $M_\phi : \mathbb{R}^e \rightarrow \mathbb{R}^d$ be a linear model. Then, even if θ and ϕ are such that $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$, this does not provide any guarantees on l_-, l_+ . When furthermore $\dim(\ker(M_\phi)) > 0$ (in particular for $d > e$), a vanishing \mathcal{L}_{mse} does not even guarantee E_θ to be a bilipschitz representation at all.

Prop. 1 shows that learning to predict the expert features does not provide any guarantees for the two properties (P1) or (P2). The proof can be found in App. C.1.

3.4. Contrastive Learning with Continuous Features

The preceding section shows that two straightforward approaches do not guarantee an encoder E to have the desired properties. We thus aim to find a new objective for CL which ensures the properties (P1) and (P2) to be fulfilled, and to a good degree at that. For this, we first propose to generalize the discrete similarity measure used within Eq. 1 to continuous labels or features. We do this by defining

$$s_{ij} := 1 - \frac{\|f_i - f_j\|_2}{\max_{k,l} \|f_k - f_l\|_2}, \quad (2)$$

where the maximum can either be taken over the complete dataset X or only over a subsample (batch). Other ways of defining the similarity measure are equally possible, see also the possibilities discussed in Sec. 3.6. We point out that, for discrete class labels which we assume here to be one-hot-encoded, our generalized similarity measure reduces to the discrete similarity measure $s_{ij} = \delta_{f_i=f_j}$ from Sec. 3.1. While this is not a necessary condition, it allows us to treat discrete and continuous features f in a uniform manner with our generalized similarity measure. When plugging the new similarity (Eq. 2) into the original pair-loss (Eq. 1), the resulting loss function encourages the desired properties (P1) and (P2) to be fulfilled, as shown below. However, the resulting loss function has the issue of discontinuities in its derivative, similar to versions of the pair-loss (Eq. 1) for continuous features (see also App. E), which might cause instabilities during optimization. We remedy this by designing a novel version of the continuous pair-loss, which we name the *quadratic contrastive loss*:

$$\mathcal{L}_{quad}(E(X), F) := \frac{1}{N^2} \sum_{i,j=1}^N ((1 - s_{ij})\Delta - D_{ij})^2, \quad (3)$$

where again $D_{ij} := \|E(x_i) - E(x_j)\|_2$. \mathcal{L}_{quad} has the same minimum as the pair-loss (Eq. 1), but possesses continuous derivatives w.r.t. D_{ij} and thus E . Furthermore, in contrast to the procedure from Prop. 1, minimizing \mathcal{L}_{quad} does lead to (optimal) guarantees for l_- and l_+ :

Proposition 2. Let \mathcal{L}_{quad} be the quadratic contrastive loss (3) and E_θ the encoder. If θ is such that $\mathcal{L}_{quad}(E_\theta(X), F) = 0$, then E_θ is a $[l_-, l_+]$ -bilipschitz representation with $l_- = l_+ = \max_{i,j} \|f_i - f_j\|_2 / \Delta$

While Prop. 2 gives guarantees for the desired properties only on the training set X (App. C.2), this can be boosted to statistical bounds for (P1) and (P2) on unseen test data (App. D). Such guarantees are relevant for downstream tasks such as outlier detection, searching, or (safe) active learning.

3.5. Implicit Hard-Negative Mining

In the previous section we introduced the quadratic contrastive loss \mathcal{L}_{quad} (Eq. 3), which puts equal weight on all pairs of datapoints. While this works well, the performance in practice often improves when higher weight is put on high-loss datapoint pairs. Such a strategy is known as “hard-negative mining” in CL (Wang & Liu, 2021) and control-theory (Busseti et al., 2016), and also improves our method (Sec. 4.6). Rather than explicitly selecting the high-loss pairs, we perform *implicit* hard-negative mining through a version of the softmax-function. This changes the loss to:

$$\mathcal{L}_{ExpCLR}^\tau(E(X), F) = \tau \log \left[\sum_{i,j=1}^N \frac{\exp\left(\frac{L_{ij}}{\tau}\right)}{N^2} \right], \quad (4)$$

where $L_{ij} := ((1 - s_{ij})\Delta - D_{ij})^2$ and $\tau \in \mathbb{R}^+$ is a temperature hyperparameter. The following proposition demonstrates that by changing τ , one can control the strength of the implicit hard-negative mining:

Proposition 3.

- (a) In the limit $\tau \rightarrow 0$, minimizing $\mathcal{L}_{ExpCLR}^\tau(E(X), F)$ is equivalent to minimizing $\mathcal{L}_{max}(E(X), F) := \max_{i,j} L_{ij}$.
- (b) In the limit $\tau \rightarrow \infty$, minimizing $\mathcal{L}_{ExpCLR}^\tau(E(X), F)$ is equivalent to minimizing $\mathcal{L}_{quad}(E(X), F)$.

Thus, the amount of hard-negative mining one wants to apply can be controlled by τ (proof in App. C.3). Similar limits exist for the NCE-Loss (Wang & Liu, 2021). A gradient-level analysis of the hard-negative mining loss can be found in App. F. We refer to the method of minimizing the loss \mathcal{L}_{ExpCLR} (Eq. 4), i.e. our quadratic contrastive loss with implicit hard-negative mining, as *ExpCLR*.

3.6. Practical Considerations for Contrastive Learning with Continuous Expert Features

Here we discuss two practical variations of \mathcal{L}_{ExpCLR} we use in our experiments. First, while we introduced a simple similarity measure in Eq. 2, in practice we use its square,

$$s_{ij} = \left(1 - \frac{\|f_i - f_j\|_2}{\max_{k,l} \|f_k - f_l\|_2} \right)^2, \quad (5)$$

Table 1: **Dataset Information:** Number of train and test samples, sample length T , dimension of expert features d , number of signal channels c , and number of classes C .

Dataset	#Train	#Test	Length	#Exp. feat.	#Chan.	#Class.
HAR	7352	2947	128	561	9	6
SleepEDF	35503	6805	3000	29	1	5
Waveform	59922	16645	2500	176	2	4

due to its consistent superior performance in experiments. Other similarity measures, such as $s_{ij} = \exp(-\|f_i - f_j\|_2^2 / \sigma)$ (Kim et al., 2019) with a hyperparameter σ , would be equally possible. For an empirical comparison of the three similarity measures, see Sec. 4.6.

Further, instead of using the unnormalized Euclidean metric and following (Kim et al., 2021), we normalize the Euclidean distance with $\mu_i = (1/N) \sum_j \|E(x_i) - E(x_j)\|_2$, i.e. we use $D_{ij} := \|E(x_i) - E(x_j)\|_2 / \mu_i$. Alternative normalizations would be possible as well.

4. Results

4.1. Datasets and Expert Features

In the following we compare ExpCLR to several state-of-the-art methods on three real-world time-series datasets. We start by introducing the three datasets; see Tab. 1 for detailed specifications.

Human Activity Recognition (HAR): The HAR dataset (Cruciani et al., 2019) contains multi-channel sensor signals of 30 subjects, each performing one out of six possible activities. A Samsung Galaxy S2 device embeds accelerometers and gyroscopes which collected the data at a constant rate of 50Hz. In addition, the dataset already contains a 561-dimensional expert feature vector for each sample.

Sleep Stage Classification (SleepEDF): In this classification task the goal is to classify five sleep stages from single-channel EEG signals, each sampled at 100Hz. The dataset originates from (Goldberger et al., 2000; Kemp et al., 2000) and subjects are selected and preprocessed following previous studies (Eldele et al., 2021). We equip each signal with expert features computed from the time and frequency domain, as suggested and identified in (Huang et al., 2020).

MIT-BIH Atrial Fibrillation (Waveform): This dataset (Goldberger et al., 2000) contains 23 long-term ECG recordings of humans suffering from atrial fibrillation. Two ECG signals are sampled at a constant rate of 250Hz and distinguish four different classes. We utilize the expert features designed by (Goodfellow et al., 2017) specifically for the atrial fibrillation classification task.

To evaluate the expert features we report their resulting linear and KNN ($k = 1$) classification accuracies in Tab. 2.

Table 2: **Unsupervised Learning Comparison:** Comparison of ExpCLR to state-of-the-art unsupervised representation learning methods on the HAR, SleepEDF and Waveform datasets. The table shows the mean performance and standard error over five independent trials for the linear and KNN ($k = 1$) classification accuracies. For better comparison we also include the performance of the representations obtained from supervised learning (full labeled data), of the expert features, and of the randomly initialized encoder network. Overall, ExpCLR outperforms the other unsupervised methods and even surpasses the supervised performance on the SleepEDF dataset.

Dataset	HAR		SleepEDF		Waveform	
	Performance (in %)	Lin. Acc.	KNN Acc.	Lin. Acc.	KNN Acc.	Lin. Acc
Cross-Entropy (S)	96.47 +/- 0.09	96.57 +/- 0.09	80.90 +/- 0.13	80.80 +/- 0.17	97.03 +/- 0.09	96.97 +/- 0.11
Expert Features	96.01 +/- 0.00	87.90 +/- 0.00	77.00 +/- 0.00	73.50 +/- 0.00	44.40 +/- 0.00	92.00 +/- 0.00
Random Init	67.74 +/- 0.59	75.02 +/- 0.48	55.14 +/- 0.46	43.78 +/- 0.28	54.56 +/- 1.04	55.20 +/- 0.62
ExpCLR (U)	91.18 +/- 0.41	88.72 +/- 0.22	81.84 +/- 0.12	74.82 +/- 0.12	92.64 +/- 0.88	88.30 +/- 0.98
SimCLR (U)	90.70 +/- 0.30	88.94 +/- 0.46	68.32 +/- 0.16	46.28 +/- 0.38	62.28 +/- 4.76	76.58 +/- 0.96
SleepPriorCL (U)	88.98 +/- 0.25	83.50 +/- 0.23	78.56 +/- 0.05	71.68 +/- 0.07	92.06 +/- 0.36	88.70 +/- 0.64
Kim et al. (2021) (U)	89.02 +/- 0.15	86.72 +/- 0.31	75.70 +/- 0.25	61.12 +/- 0.27	83.96 +/- 1.75	81.80 +/- 0.93
TS-TCC (U)	90.57 +/- 0.15	–	80.68 +/- 0.24	–	82.17 +/- 2.53	–
TREBA Contrastive Loss (U)	78.40 +/- 1.79	65.90 +/- 0.16	77.73 +/- 0.63	70.20 +/- 0.08	90.53 +/- 0.68	81.13 +/- 0.52
Expert Feature Decoding (U)	85.20 +/- 2.69	79.83 +/- 3.01	80.73 +/- 0.05	74.97 +/- 0.45	91.83 +/- 1.10	82.70 +/- 3.80

4.2. Implementation Details – Model and Training

Next, we briefly present our model architecture and how we train our models in each setting; more details can be found in App. B.3. To capture relevant temporal properties and to improve training stability (Bai et al., 2018), we choose as a base encoder temporal convolutional network (TCN) (Lea et al., 2017) layers in a ResNet (He et al., 2016) architecture with eight such temporal blocks. Note that ExpCLR is not restricted to this architecture. To reach a pre-defined embedding dimensionality we add a two-layer fully connected neural network on top of the ResNet base encoder to arrive at our backbone encoder network.

In our work we consider three different modes of training:

1. The *unsupervised* (U) training mode, where the encoder is optimized with the respective contrastive loss function on the input data time-series X and the expert features F only.
2. The *supervised* (S) training mode, where the encoder is trained with either a supervised contrastive loss, which uses labels y_i instead of expert features f_i , or with the cross-entropy loss function on the input time-series X and the whole or part of the labels Y .
3. The *semi-supervised* (SS) training mode, where the encoder is first trained with the unsupervised training mode on the whole training set (X, F) and then this pretrained encoder is fine-tuned with a supervised training step on some percentage of the labels Y .

While for hyperparameter optimization we split the training set X into 80% training and 20% validation data, for our

comparisons experiments we make use of the full training set and evaluate the representations on the test set. The number of epochs for each dataset is selected such that all algorithms are able converge. For the optimization step we used the Adam optimizer with parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$ and exponential decay $\gamma = 0.99$. To enable a fair comparison between ExpCLR and the competing methods, we optimize the learning rate for each method and dataset individually via a grid search and identify $\tau = 1$, $\Delta = 1$ (Eq. 4), embedding dimension $e = 100$ and batch size of 64 as a good compromise over all datasets and algorithms. For more information on model- and loss-specific parameters, see Sec. 4.6, App. A.2 and App. B.3.

To verify the goodness of our representations, we evaluate two kinds of classifiers on the representation: Linear classifiers perform well for representations where all classes are linearly separable. Second, KNN classifiers can even perform well when classes are not linearly separable, but tend to perform worse for clusters that are not separated by a large margin; this problem is most apparent for small k . We thus use the performance difference of the linear and KNN ($k = 1$) classifier to investigate how well different classes are mapped into individual well-separated clusters.

4.3. Competing Methods

For the unsupervised comparison of ExpCLR to other methods, we distinguish two different groups of CL algorithms. The first group consist of algorithms, which use transformations of the input data to create positive samples. Here we compare to SimCLR (Chen et al., 2020) and further to TS-TCC (Eldele et al., 2021), that combines classical CL with time-series forecasting to learn representations.

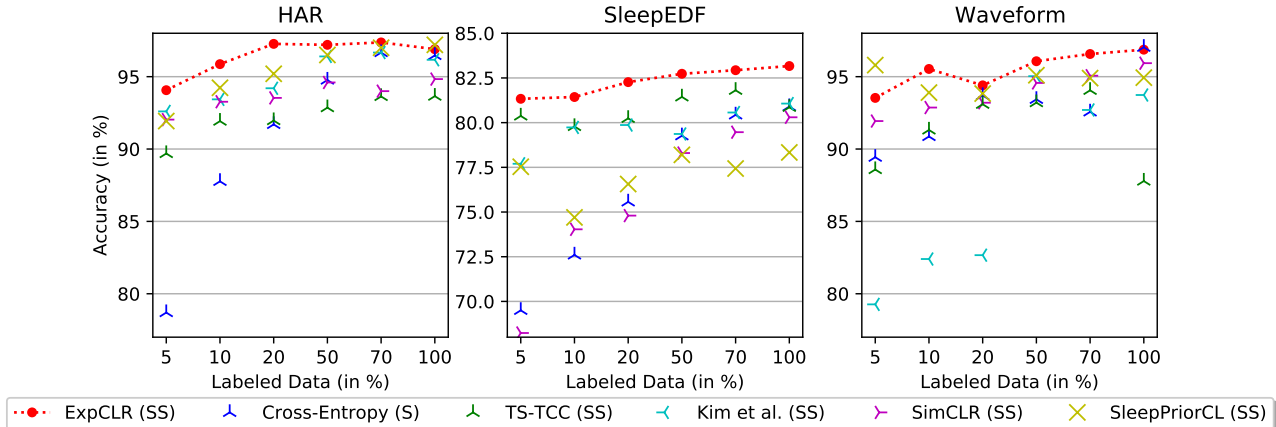


Figure 1. **Semi-Supervised Learning Comparison:** Comparison of ExpCLR (SS) to semi-supervised representation learning methods on the HAR, SleepEDF, and Waveform datasets. Shown are the mean linear classification accuracies for different percentages of labeled data available during the supervised fine-tuning step. Across all datasets, ExpCLR outperforms all other methods. Further, while ExpCLR achieves consistent performance on all three datasets, all competing methods have drastically varying performance across datasets.

For SimCLR we tested a range of different augmentations and found scaling and dropout to work the best, while for TS-TCC we employ weak and strong augmentations. Further details can be found in App. B.4. The second group includes methods which also use expert features. SleepPriorCL (Zhang et al., 2021) and the contrastive loss used by TREBA (Sun et al., 2021) both create expert feature based pseudo-labels via some form of discretization, which are then used with a version of the supervised contrastive loss introduced by (Khosla et al., 2020). Another method in this group is introduced by (Kim et al., 2021), a state-of-the-art metric learning method. It aims to achieve the same goal as we do and try to pull similar points closer together (Kim et al., 2021), while pushing dissimilar ones further apart. We simply replace their teacher model output with our expert features to be comparable. Lastly, we also compare to the embedding learned by Expert Feature Decoding: Here, the embedding is given by the output of the penultimate layer of a network that is trained by learning to predict the expert features from the input time-series (Sun et al., 2021). During training we minimized the MSE-loss and add a projection layer to the architecture used by the other methods. Expert Feature Decoding is used as part of the TREBA-objective (Sun et al., 2021) and is discussed theoretically in Sec. 3.3.

For the supervised fine-tuning step we use the approach described in the original works or utilize the natural extension for each algorithm. For SimCLR we use supervised CL (Khosla et al., 2020) and for (Kim et al., 2021) we use the pair-Loss (Hadsell et al., 2006; Le-Khac et al., 2020), because this is the loss that it naturally reduces to for class labels. Further, as ExpCLR allows to simply replace expert features with labels in order to perform a supervised fine-tuning step or a fully supervised training, we do this.

We selected the competing methods to cover a broad range of algorithms. All methods can be considered state-of-the-art in their respective domains. We implemented all methods except for TS-TCC inside our repository.

4.4. Comparison for Unsupervised Representation Learning

In this section we compare the performance of ExpCLR against several state-of-the-art unsupervised representation learning methods, using a linear and a KNN ($k = 1$) classifier on top of the learned embedding. We compare ExpCLR on all three datasets against SimCLR, TS-TCC, SleepPriorCL, TREBA Contrastive Loss, Expert Feature Decoding and (Kim et al., 2021). In addition, we use the randomly initialized encoder network performance (Random Init), an encoder trained with a supervised cross-entropy loss, using the full labeled dataset, and the performance on the expert features themselves as baseline comparisons. The results of the comparison are shown in Tab. 2.

The superior performance of ExpCLR can be clearly seen, as we only perform on par with SimCLR on HAR and with SleepPriorCL on Waveform w.r.t. KNN ($k = 1$) accuracy. ExpCLR also shows much higher consistency across all datasets, while most other algorithms significantly underperform on at least one of the datasets. Further, we are even able to exceed the expert feature performance on at least one performance metric on all three datasets. This could indicate that ExpCLR is able to learn new additional features from the raw time-series data on top of the provided expert features. In addition, we can even surpass the supervised performance on the SleepEDF dataset. This underlines how powerful the approach of ExpCLR.

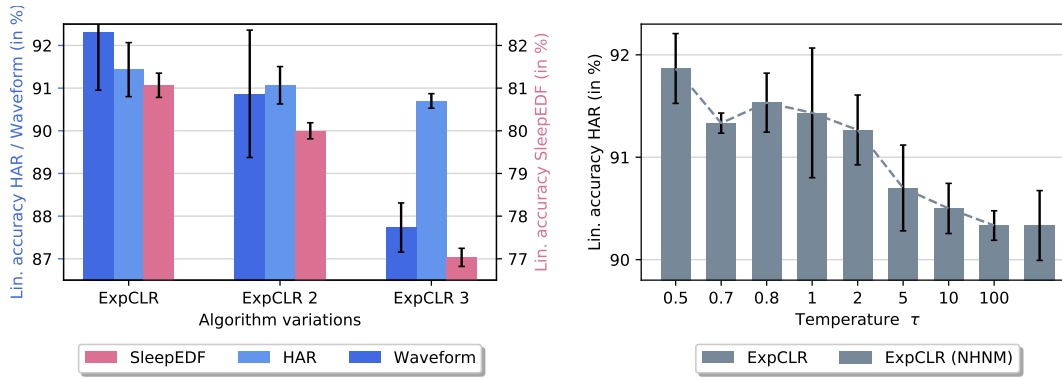


Figure 2. **Ablation Experiments:** The left panel shows a comparison of three different similarity measures in ExpCLR in the unsupervised training mode. While ExpCLR uses the similarity measure (Eq. 5), ExpCLR 2 uses the similarity (Eq. 2), and ExpCLR 3 uses the exponential similarity introduced by (Kim et al., 2021) (Sec. 3.6). Our choice clearly outperforms the other ones. The right panel investigates the effectiveness of our implicit hard-negative mining scheme (Sec. 3.5) for the unsupervised setting on the HAR dataset. We plot the linear accuracies over different values of τ and also compare these to a version of ExpCLR without hard-negative mining (NHNM), i.e. the quadratic contrastive loss (Eq. 3). The results show that hard-negative mining can improve the performance significantly.

4.5. Comparison for Semi-Supervised Representation Learning

Here, we compare ExpCLR to the competing methods in the semi-supervised setting. In this setting we fine-tune a representation, learned in the unsupervised setting, using some fraction of the full labeled data. The results of the comparison for the HAR, SleepEDF and Waveform datasets for labeled data ratios of 5%, 10%, 20%, 50%, 70% and 100% are shown in Fig. 1 (see App. A.1 for the KNN accuracy).

ExpCLR consistently surpasses the accuracy of all competing methods across all datasets over different label percentages. Further, using only 20% labeled data ExpCLR (SS) is able to outperform any competing method with any amount of labeled data, even for 100%, on the HAR and SleepEDF datasets. In addition, ExpCLR’s unsupervised pretraining increases the label efficiency significantly, since ExpCLR is able to sustain its performance attained on the fully labeled dataset up to a minor decrease in accuracy of less than 3.5% for the lowest labeled data percentage. In contrast, the representations learned by supervised cross-entropy drop by more than 17.5%.

The consistent superior performance of ExpCLR across all datasets is noteworthy as it does not use any data transformations. Further, the comparison clearly shows the superior performance of our loss function compared to Kim et al. (2021). Another interesting observation is that almost all methods that make use of an unsupervised pretraining outperform supervised methods for low percentages of labeled data. This underlines the importance and effectiveness of finding good approaches that make use of unlabeled data, as ours. All numerical values of Fig. 1 incl. variances can be found in App. A.5 (see also App. A.6).

4.6. Ablation Studies

To guide the design of our ExpCLR method we conducted several ablation studies. The first one investigates the choice of similarity measure in ExpCLR and is shown in the left panel of Fig. 2. We compare both similarities (Eq. 2) and (Eq. 5) from our methods section and the Gaussian similarity (Sec. 3.6) introduced by Kim et al. (2021) in the unsupervised setting for all three real-world datasets. The results show that the quadratic similarity measure (Eq. 5) outperforms the other similarity measures, therefore justifying our choice.

The second ablation study investigates the effectiveness of the hard-negative mining strategies introduced in Sec. 3.5, which ExpCLR employs. The right panel of Fig. 2 shows the unsupervised accuracy of ExpCLR for several values of the temperature τ and also compares to ExpCLR without hard-negative mining (NHNM). Using the insights gained from Prop. 3 (a), Fig. 2 shows that stronger hard-negative mining (decreasing τ) can improve the performance. Further, as shown theoretically in Prop. 3 (b), the results nicely demonstrate that the performance converges to ExpCLR (NHNM) for large values of τ . More ablation studies and a sensitivity analysis for the hyperparameter Δ , the batch-size and the dimension of the embedding can be found in the appendix (see App. A.2).

5. Discussion

In this paper we introduce ExpCLR, a novel contrastive representation learning algorithm that can utilize continuous or discrete expert features. We first propose two properties a useful time-series representation should fulfill. In a second step, we design ExpCLR to be applicable in the unsuper-

vised and semi-supervised domains and show that the loss function we devised for ExpCLR leads to a representation that encourages both properties. We demonstrate on an array of experiments the superior performance of ExpCLR compared to state-of-the-art methods, sometimes exceeding their accuracies despite using only a fraction of the labels.

We see in ExpCLR an alternative to the classical transformation-based contrastive learning (CL) approaches, as ExpCLR does not make use of any transformations. Nevertheless, ExpCLR is able to outperform these classical approaches in both the un- and semi-supervised settings. This is especially noteworthy for datasets where domain experts can provide valuable expert features. Thus, ExpCLR is applicable to any dataset for which expert features are available, and is not limited to time-series datasets. In addition, ExpCLR can also be applied to supervised CL with datasets containing continuous labels, e.g. regression tasks such as pose estimation.

We envision our ExpCLR approach to not only serve as a standalone method for representation learning, but to be applicable to any task or dataset where (continuous) expert features are available in order to infuse expert knowledge into neural networks. Apart from pretraining, this might be achieved by jointly training our ExpCLR loss together with a task-specific loss function, as done by TREBA (Sun et al., 2021), which should increase the task performance while being more label-efficient.

References

- Bai, S., Kolter, J. Z., and Koltun, V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.
- Bromley, J., Bentz, J. W., Bottou, L., Guyon, I., LeCun, Y., Moore, C., Säckinger, E., and Shah, R. Signature verification using a “siamese” time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993.
- Busseti, E., Ryu, E. K., and Boyd, S. Risk-constrained Kelly gambling. *The Journal of Investing*, 25(3):118–134, 2016.
- Chattha, M. A., Siddiqui, S. A., Malik, M. I., van Elst, L., Dengel, A., and Ahmed, S. KINN: Incorporating expert knowledge in neural networks. *arXiv preprint arXiv:1902.05653*, 2019.
- Chechik, G., Sharma, V., Shalit, U., and Bengio, S. Large scale online learning of image similarity through ranking. *Journal of Machine Learning Research*, 11(3), 2010.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.
- Chuang, C.-Y., Robinson, J., Yen-Chen, L., Torralba, A., and Jegelka, S. Debiased contrastive learning. *arXiv preprint arXiv:2007.00224*, 2020.
- Cruciani, F., Sun, C., Zhang, S., Nugent, C., Li, C., Song, S., Cheng, C., Cleland, I., and Mccullagh, P. A public domain dataset for human activity recognition in free-living conditions. In *2019 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/IOP/SCI)*, pp. 166–171. IEEE, 2019.
- Eldele, E., Ragab, M., Chen, Z., Wu, M., Kwok, C. K., Li, X., and Guan, C. Time-series representation learning via temporal and contextual contrasting. *arXiv preprint arXiv:2106.14112*, 2021.
- Franceschi, J.-Y., Dieuleveut, A., and Jaggi, M. Unsupervised scalable representation learning for multivariate time series. *arXiv preprint arXiv:1901.10738*, 2019.
- Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.
- Goodfellow, S. D., Goodwin, A., Greer, R., Laussen, P. C., Mazwi, M., and Eytan, D. Classification of atrial fibrillation using multidisciplinary features and gradient boosting. In *2017 Computing in Cardiology (CinC)*, pp. 1–4. IEEE, 2017.
- Gou, J., Yu, B., Maybank, S. J., and Tao, D. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- 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., et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 297–304. JMLR Workshop and Conference Proceedings, 2010.

- Hadsell, R., Chopra, S., and LeCun, Y. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pp. 1735–1742. IEEE, 2006.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Hu, Z., Ma, X., Liu, Z., Hovy, E., and Xing, E. Harnessing deep neural networks with logic rules. *arXiv preprint arXiv:1603.06318*, 2016.
- Huang, W., Guo, B., Shen, Y., Tang, X., Zhang, T., Li, D., and Jiang, Z. Sleep staging algorithm based on multichannel data adding and multifeature screening. *Computer Methods and Programs in Biomedicine*, 187:105253, 2020. ISSN 0169-2607. doi: <https://doi.org/10.1016/j.cmpb.2019.105253>. URL <https://www.sciencedirect.com/science/article/pii/S0169260719304602>.
- Iwana, B. K. and Uchida, S. An empirical survey of data augmentation for time series classification with neural networks. *Plos one*, 16(7):e0254841, 2021.
- Kallidromitis, K., Gudovskiy, D., Kazuki, K., Iku, O., and Rigazio, L. Contrastive neural processes for self-supervised learning. In *Asian Conference on Machine Learning*, pp. 594–609. PMLR, 2021.
- Kemp, B., Zwinderman, A., Tuk, B., Kamphuisen, H., and Oberye, J. Analysis of a sleep-dependent neuronal feedback loop: the slow-wave microcontinuity of the eeg. *IEEE Transactions on Biomedical Engineering*, 47(9): 1185–1194, 2000. doi: 10.1109/10.867928.
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., and Krishnan, D. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020.
- Kim, S., Seo, M., Laptev, I., Cho, M., and Kwak, S. Deep metric learning beyond binary supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2288–2297, 2019.
- Kim, S., Kim, D., Cho, M., and Kwak, S. Proxy anchor loss for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3238–3247, 2020.
- Kim, S., Kim, D., Cho, M., and Kwak, S. Embedding transfer with label relaxation for improved metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3967–3976, 2021.
- Le-Khac, P. H., Healy, G., and Smeaton, A. F. Contrastive representation learning: A framework and review. *IEEE Access*, 2020.
- Lea, C., Flynn, M. D., Vidal, R., Reiter, A., and Hager, G. D. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 156–165, 2017.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Park, W., Kim, D., Lu, Y., and Cho, M. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3967–3976, 2019.
- Sun, J. J., Kennedy, A., Zhan, E., Anderson, D. J., Yue, Y., and Perona, P. Task programming: Learning data efficient behavior representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2876–2885, 2021.
- Tamkin, A., Wu, M., and Goodman, N. Viewmaker networks: Learning views for unsupervised representation learning. *arXiv preprint arXiv:2010.07432*, 2020.
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., and Isola, P. What makes for good views for contrastive learning? *arXiv preprint arXiv:2005.10243*, 2020.
- Tonekaboni, S., Eytan, D., and Goldenberg, A. Unsupervised representation learning for time series with temporal neighborhood coding. In *International Conference on Learning Representations*, 2020.
- Wang, F. and Liu, H. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2495–2504, 2021.
- Wang, T. and Isola, P. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pp. 9929–9939. PMLR, 2020.
- Yu, L., Yazici, V. O., Liu, X., Weijer, J. v. d., Cheng, Y., and Ramisa, A. Learning metrics from teachers: Compact networks for image embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2907–2916, 2019.

Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. Barlow twins: Self-supervised learning via redundancy reduction. *arXiv preprint arXiv:2103.03230*, 2021.

Zhang, H., Wang, J., Xiao, Q., Deng, J., and Lin, Y. Sleep-PriorCL: Contrastive representation learning with prior knowledge-based positive mining and adaptive temperature for sleep staging. *arXiv preprint arXiv:2110.09966*, 2021.

Zheng, M., Wang, F., You, S., Qian, C., Zhang, C., Wang, X., and Xu, C. Weakly supervised contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10042–10051, 2021.

Supplementary Material for Utilizing Expert Features for Contrastive Learning of Time-Series Representations

A. Additional Experiments

A.1. Figure with KNN ($k = 1$) Classification Results for the Semi-Supervised Experiments

See Fig. 3 for a comparison of the KNN ($k = 1$) accuracies on our three real-world datasets (analogous to Fig. 1 for the linear classification accuracy). For all numerical values including standard deviations, see App. A.6.

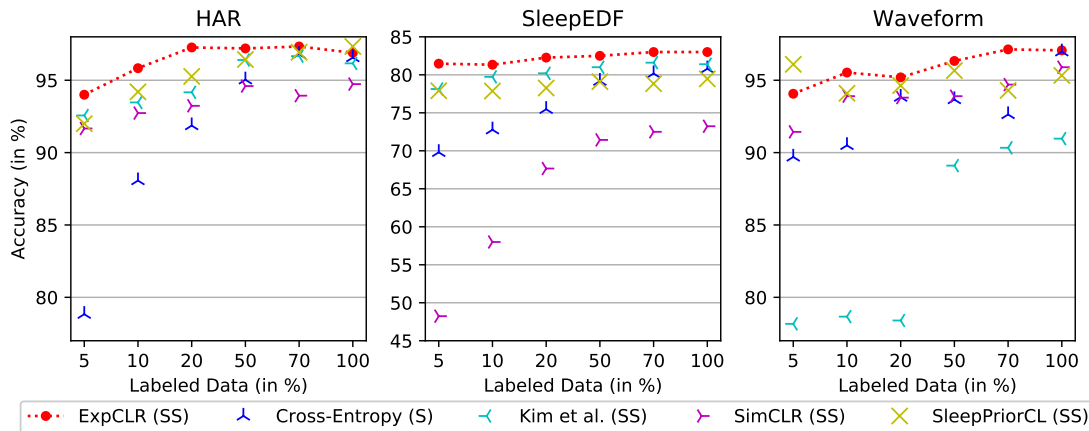


Figure 3. **Semi-Supervised Comparison KNN ($k = 1$):** The figure shows the same setting as Fig. 1, but instead of the linear classification accuracy all methods are evaluated with the KNN ($k = 1$) classification accuracy.

A.2. Sensitivity Analysis

We perform sensitivity analysis for Δ from Eq. 3, batch size during training, and dimension of e the embedding space to which the input samples are mapped. To analyse the sensitivity, we train ExpCLR in unsupervised training mode and evaluate the linear accuracy of the embedding space. The results are shown in Fig. 4. One can see that:

- **Left plot:** There is an optimal value for Δ which is robust against modifications within a certain range. For too low or too high Δ , linear accuracy tends to deteriorate because pushing the embedding to those margins can result in more difficult optimization leading to training instability. For all contrastive learning approaches that include this hyperparameter Δ in their losses, we choose $\Delta = 1$ for all datasets.
- **Middle plot:** For small batch sizes, performance drops as there might not be enough dissimilar samples within the same batch (Chen et al., 2020). Also for too large batch sizes the performance decreases for all datasets consistently. For all algorithms and experiments we choose the $batchsize = 64$.
- **Right plot:** Higher embedding dimensions tend to improving accuracies because the placement of embedding vectors and their separation might be easier in higher dimensional space. Note that although a low-dimensional embedding is generally preferable, too low dimensions may cause unstable training as different input vectors are forced to be pushed towards similar representations which goes against the nature of the loss function from Eq. 4. We observe an embedding size of $e = 100$ as applicable and choose it over all datasets and algorithms consistently.

A.3. Comparison ExpCLR (S) vs. ExpCLR (SS)

Figs. 5 and 6 show a comparison of the supervised (S) and semi-supervised (SS) versions of ExpCLR for the linear and KNN ($k = 1$) classification accuracies they achieve, plotted over the fraction of labeled data used.

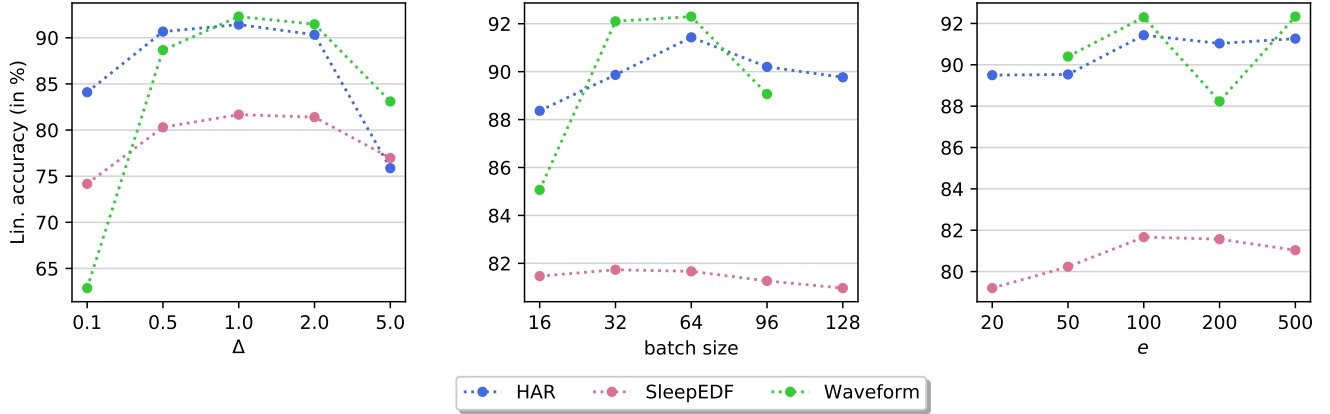


Figure 4. Sensitivity Analysis ExpCLR Unsupervised: The figure shows three sensitivity analyses for ExpCLR unsupervised training, evaluated by linear accuracy. The left panel shows results for different Δ in Eq. 4, in the middle panel for different batch sizes, and in the right panel for different embedding dimensions e .

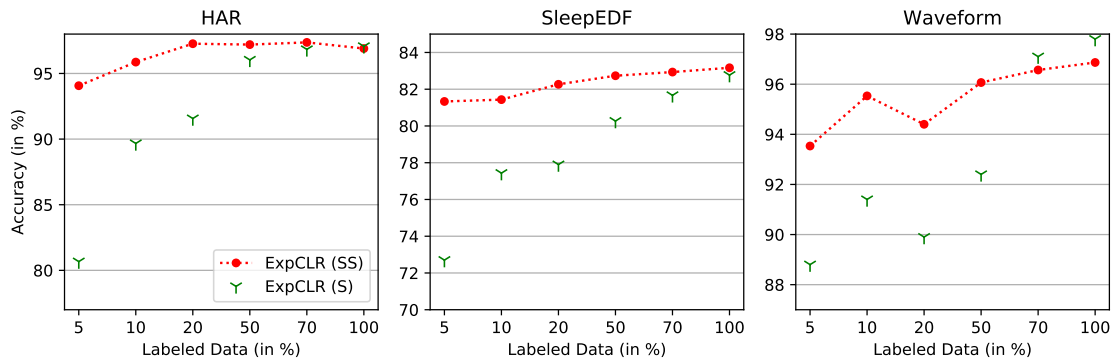


Figure 5. ExpCLR (SS) vs ExpCLR (S) on Linear Classification Accuracies: Shown is the same setting as in Fig. 1 but only comparing the ExpCLR method in the semi-supervised and supervised settings. The results show that the unsupervised pretraining enables ExpCLR (SS) to clearly surpass the performance obtained by ExpCLR (S) overall, especially for lower percentages of labeled data.

A.4. Comparison of ExpCLR (SS) vs. Competing Methods Using ExpCLR (S) for the Fine-Tuning Step

See Fig. 7 for a comparison of ExpCLR (SS) with the competing methods, where instead of using the standard fine-tuning step for each method (as was shown in Fig. 1 and described in Sec. 4.3) we apply ExpCLR (S) as the fine-tuning step for each competing method.

A.5. Tabular Data of the Linear Classification Accuracies for the Semi-Supervised Experiments

The tables Tab. 3, 4 and 5 show the tabular data for Fig. 1 (Sec. 4.5).

A.6. Tabular Data of the KNN ($k = 1$) Classification Accuracies for the Semi-Supervised Experiments

The tables Tab. 6, 7 and 8 show the tabular data for Fig. 3 (App. A.1 and Sec. 4.5).

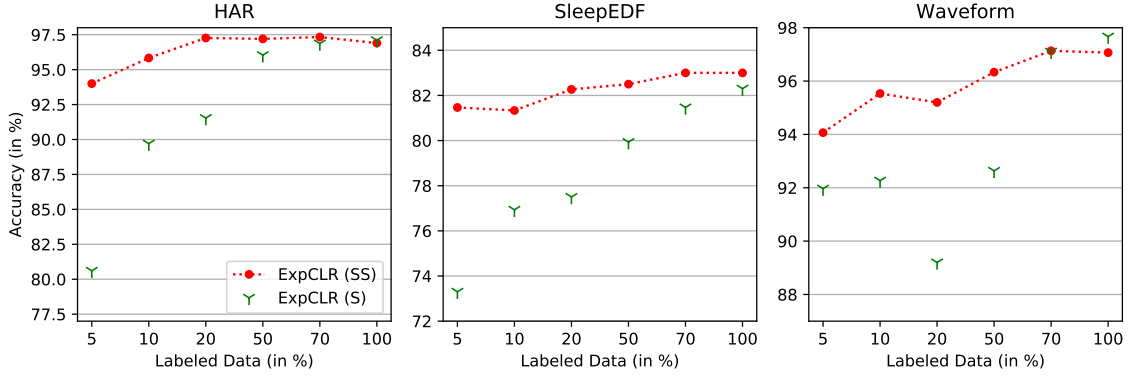


Figure 6. ExpCLR (SS) vs ExpCLR (S) on KNN ($k = 1$) Classification Accuracies: The figure shows the same comparison as shown in Fig. 5 but instead of the linear accuracies it shows uses KNN ($k = 1$) accuracies to evaluate the quality of the representations. The results complement the findings of Fig. 5.

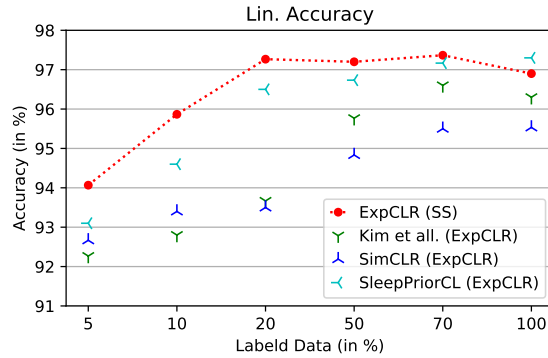


Figure 7. Comparisons of ExpCLR (SS) vs Competing Methods with ExpCLR (S) Fine-Tuning: The figure shows a comparison on the HAR dataset of ExpCLR (SS) with the competing methods but instead of using the standard fine-tuning step for each step, we apply ExpCLR (S) as a fine-tuning step for each competing method. The results indicate that while ExpCLR (S) performs well as a fine-tuning method, a key ingredient for the good performance of ExpCLR (SS) is the pretraining step with ExpCLR (U), using expert features.

B. Experimental Details

B.1. Datasets

In the following we list the sources, further information, and implementation details of the datasets and expert features we used in our work.

HAR: The HAR dataset aims to classify six different activity states of humans, namely: walking, walking upstairs, walking downstairs, sitting, standing, laying. They collected the data using a mounted Samsung Galaxy S2 device where a triaxial acceleration and gyroscope sensor is installed. We downloaded the dataset from the UCI Machine Learning Repository (<https://archive.ics.uci.edu/ml/datasets/human+activity+recognition+using+smartphones>) and preprocessed the data as (Eldele et al., 2021) did in the corresponding repository (<https://github.com/emadeldeen24/TS-TCC>).

SleepEDF: For sleep stage classification they differ five different classes: Wake (W), Non-rapid eye movement (N1, N2, N3) and Rapid Eye Movement (REM). We downloaded the dataset from the PyhsioNet database (<https://physionet.org/content/sleep-edf/1.0.0>) and load/preprocessed the data like (Eldele et al., 2021) did in the repository (<https://github.com/emadeldeen24/TS-TCC>).

Waveform: This dataset distinguishes four different classes: Atrial Fibrillation (AFIB), atrial flutter (AFL), AV junctional rhythm (J), and all other rhythms (N). We downloaded the data from the PyhsioNet database

Table 3: Tabular data for the semi-supervised comparison shown in Fig. 1 for linear classification accuracies on the HAR dataset.

Labeled Data	5%	10%	20%	50%	70%	100%
ExpCLR (SS)	94.07 +/- 0.71	95.87 +/- 0.38	97.27 +/- 0.54	97.20 +/- 0.37	97.37 +/- 0.27	96.90 +/- 0.29
ExpCLR (S)	80.67 +/- 2.65	89.67 +/- 0.31	91.57 +/- 0.83	96.03 +/- 0.26	96.83 +/- 0.15	97.07 +/- 0.49
Cross-Entropy (S)	78.70 +/- 0.90	87.77 +/- 0.19	91.70 +/- 0.50	94.77 +/- 0.58	96.67 +/- 0.10	96.47 +/- 0.12
TS-TCC (SS)	89.69 +/- 0.62	91.93 +/- 0.18	91.97 +/- 0.21	92.86 +/- 0.05	93.63 +/- 0.25	93.65 +/- 0.39
Kim et al. (2021) (SS)	92.60 +/- 0.64	93.43 +/- 0.93	94.20 +/- 0.87	96.40 +/- 0.08	96.67 +/- 0.26	96.17 +/- 0.18
SimCLR (SS)	92.03 +/- 0.59	93.27 +/- 0.41	93.53 +/- 0.27	94.60 +/- 0.08	94.00 +/- 0.29	94.83 +/- 0.15
SleepPriorCL (SS)	91.93 +/- 0.57	94.23 +/- 0.47	95.20 +/- 0.75	96.50 +/- 0.37	97.00 +/- 0.28	97.20 +/- 0.26

Table 4: Tabular data for the semi-supervised comparison shown in Fig. 1 for linear classification accuracies on the SleepEDF dataset.

Labeled Data	5%	10%	20%	50%	70%	100%
ExpCLR (SS)	81.33 +/- 0.07	81.43 +/- 0.10	82.27 +/- 0.20	82.73 +/- 0.03	82.93 +/- 0.07	83.17 +/- 0.18
ExpCLR (S)	72.70 +/- 0.34	77.43 +/- 0.21	77.90 +/- 0.52	80.27 +/- 0.24	81.67 +/- 0.35	82.77 +/- 0.05
Cross-Entropy (S)	69.50 +/- 0.33	72.60 +/- 0.71	75.57 +/- 0.07	79.27 +/- 0.30	80.43 +/- 0.05	80.90 +/- 0.17
TS-TCC (SS)	80.37 +/- 0.22	79.79 +/- 0.08	80.24 +/- 0.15	81.44 +/- 0.18	81.81 +/- 0.15	80.86 +/- 0.37
Kim et al. (2021) (SS)	77.70 +/- 0.50	79.73 +/- 0.03	79.87 +/- 0.32	79.37 +/- 0.03	80.57 +/- 0.72	81.07 +/- 0.34
SimCLR (SS)	68.23 +/- 0.05	74.03 +/- 0.30	74.80 +/- 0.42	78.30 +/- 0.26	79.47 +/- 0.24	80.30 +/- 0.14
SleepPriorCL (SS)	77.53 +/- 0.12	74.70 +/- 0.70	76.57 +/- 0.67	78.20 +/- 0.22	77.43 +/- 0.43	78.33 +/- 0.30

(<https://physionet.org/content/afdb/1.0.0>) and preprocessed it like (Tonekaboni et al., 2020) did in the corresponding repository (https://github.com/sanatonk/TNC_representation_learning).

B.2. Expert Features:

HAR: To obtain expert features, (Cruciani et al., 2019) first filtered the raw time-series signals in order to reduce noise. Subsequently, for selected signals they compute signal magnitudes and applied a Fast Fourier Transform. In order to equip all the resulting signals with expert features, they compute scalar attributes like maximal value, minimal value, means, energy or standard deviation for each signal. Finally they end up with a 561-dimensional expert feature vector for each sample.

SleepEDF: Following (Huang et al., 2020) we select their proposed methods to calculate expert features for sleep stage classification from ECG signals. They identify 30 suitable features from time and frequency domain for this classification task which we implemented in our repository. We couldn't find implementation details of the fractal dimension and left this feature concluding with a 29 dimensional expert feature vector.

Waveform: For this dataset (Goodfellow et al., 2017) found representative features for classification of ECG signals. They distinguish three feature types: Full waveform features which are extracted from the wavelet transformation, template features which identify medical properties of the ECG signal and mainly separate between normal rhythm and atrial fibrillation, and lastly RRI features which identify properties of important signal peaks. For our model we used the full waveform features and the RRI features as the template features are not directly implemented in (Goodfellow et al., 2017) repository (<https://github.com/Seb-Good/ecg-features/blob/f9a4c986f8e460a081c71b8e2c7e3ddb26eabae8>). Finally we get a 176 dimensional feature vector for each sample.

B.3. Implementation Details

Model: We start from TCN implementation (Bai et al., 2018) and fit the structure in order to remove time causality and same sequence length output, as we generally want a lower dimensional embedding. We apply a constant channel- and kernel size over all convolutional layers. We only increase the stride downsampling for longer sample lengths as they occur in SleepEDF and Waveform dataset.

Table 5: Tabular data for the semi-supervised comparison shown in Fig. 1 for linear classification accuracies on the Waveform dataset.

Labeled Data	5%	10%	20%	50%	70%	100%
ExpCLR (SS)	93.53 +/- 0.45	95.53 +/- 0.88	94.40 +/- 1.94	96.07 +/- 1.07	96.57 +/- 0.71	96.87 +/- 0.64
ExpCLR (S)	88.80 +/- 3.87	91.40 +/- 2.87	89.90 +/- 1.14	92.40 +/- 1.34	97.10 +/- 0.45	97.80 +/- 0.08
Cross-Entropy (S)	89.43 +/- 1.95	90.87 +/- 1.20	93.70 +/- 0.54	93.43 +/- 0.79	92.57 +/- 2.97	97.03 +/- 0.12
TS-TCC (SS)	88.58 +/- 2.18	91.32 +/- 0.58	93.09 +/- 0.40	93.21 +/- 0.87	94.06 +/- 0.30	87.79 +/- 1.52
Kim et al. (2021) (SS)	79.27 +/- 5.28	82.40 +/- 6.53	82.67 +/- 4.23	95.03 +/- 1.19	92.70 +/- 2.42	93.73 +/- 1.58
SimCLR (SS)	91.93 +/- 0.76	92.87 +/- 0.84	93.20 +/- 1.07	94.57 +/- 0.73	95.07 +/- 1.05	95.93 +/- 0.97
SleepPriorCL (SS)	95.80 +/- 0.87	93.90 +/- 1.41	93.83 +/- 0.90	95.10 +/- 0.86	94.90 +/- 0.99	94.93 +/- 0.83

Table 6: Tabular data for the semi-supervised comparison shown in Fig. 1 for KNN ($k = 1$) classification accuracies on the HAR dataset.

Labeled Data	5%	10%	20%	50%	70%	100%
ExpCLR (SS)	94.00 +/- 0.70	95.83 +/- 0.38	97.27 +/- 0.54	97.20 +/- 0.37	97.33 +/- 0.24	96.90 +/- 0.29
ExpCLR (S)	80.60 +/- 2.58	89.70 +/- 0.31	91.53 +/- 0.87	96.03 +/- 0.26	96.87 +/- 0.14	97.10 +/- 0.46
Cross-Entropy (S)	78.83 +/- 0.94	88.07 +/- 0.17	91.87 +/- 0.57	95.03 +/- 0.58	96.80 +/- 0.08	96.57 +/- 0.12
Kim et al. (2021) (SS)	92.57 +/- 0.62	93.47 +/- 0.96	94.17 +/- 0.88	96.40 +/- 0.08	96.67 +/- 0.26	96.17 +/- 0.18
SimCLR (SS)	91.67 +/- 0.56	92.73 +/- 0.36	93.23 +/- 0.34	94.60 +/- 0.12	93.93 +/- 0.30	94.73 +/- 0.03
SleepPriorCL (SS)	92.00 +/- 0.64	94.20 +/- 0.50	95.27 +/- 0.73	96.43 +/- 0.42	96.93 +/- 0.28	97.33 +/- 0.35

Training: For both, ExpCLR and the competing methods we choose individual learning rates for unsupervised and supervised training for each dataset optimized in a range of $lr \in \{5e-5, 1e-4, 5e-4, 1e-3, 3e-3, 5e-3, 7e-3, 1e-2\}$ each. As performance indicator we choose linear accuracy as we did in Sec. 4.2. The final learning rates for each algorithm for supervised (S) and unsupervised (U) training modes and each dataset are shown in Tab. 9. Further we investigate different values for the parameters τ and Δ from Eq. 4. To avoid overfitting ExpCLR w.r.t our competing methods we set τ and Δ to be the same for all datasets. Regarding training stability and accuracy we identify $\tau = 1$ and $\Delta = 1$ as a good choice.

B.4. Competing Methods:

For all methods except TS-TCC, we implemented the methods in our repository as explained in Sec. 4.3. For the comparison we only substitute the loss function used during optimization and keep the architecture and evaluation methods the same for all competing methods. As explained in the previous section, we select the learning rate based on a grid-search and choose the best learning rate individually for each combination of algorithm and dataset. Apart from this, we did not conduct an extensive search of all the respective hyperparameters but rather chose natural ones or used the parameters as described in the original works. This also applies to ExpCLR as can be seen in Fig. 2, where other $\tau < 1.0$ could improve the performance slightly. For the data transformations we tried out different combination and found dropout and scaling to work well for the respective datasets. These data transformations were then used for SimCLR.

For TS-TCC we made use of the already available repository (<https://github.com/emadeldeen24/TS-TCC>). We adapted the repository to include the exact same datasets and seeds we used and added our Waveform dataset to the repository. Further, we tried replacing the encoder used by TS-TCC with our encoder architecture, but found the performance to drop significantly. Therefore, we chose to keep the original encoder architecture.

C. Proofs

C.1. Proof of Proposition 1

Let us first consider the case $\dim(\ker(M_\phi)) = 0$. Further, we assume the input points to be generic and thus also the embedding vectors $E_i := E_\theta(x_i)$ can be considered generic. Since our goal is to show that the condition $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$ is insufficient to derive any bound on l_- and l_+ , we assume that there exist fixed $B_-, B_+ \in \mathbb{R}_+$ such that

Table 7: Tabular data for the semi-supervised comparison shown in Fig. 1 for KNN ($k = 1$) classification accuracies on the SleepEDF dataset.

Labeled Data	5%	10%	20%	50%	70%	100%
ExpCLR (SS)	81.47 +/- 0.18	81.33 +/- 0.05	82.27 +/- 0.23	82.50 +/- 0.17	83.00 +/- 0.19	83.00 +/- 0.25
ExpCLR (S)	73.30 +/- 0.17	76.93 +/- 0.44	77.50 +/- 0.63	79.93 +/- 0.12	81.47 +/- 0.40	82.30 +/- 0.05
Cross-Entropy (S)	69.80 +/- 0.54	72.77 +/- 0.67	75.47 +/- 0.20	79.17 +/- 0.26	80.17 +/- 0.10	80.80 +/- 0.22
Kim et al. (2021) (SS)	78.13 +/- 0.49	79.73 +/- 0.10	80.20 +/- 0.14	81.00 +/- 0.21	81.60 +/- 0.12	81.10 +/- 0.34
SimCLR (SS)	48.23 +/- 0.47	58.00 +/- 0.17	67.67 +/- 0.82	71.43 +/- 0.22	72.50 +/- 0.24	76.30 +/- 0.16
SleepPriorCL (SS)	77.90 +/- 0.31	77.87 +/- 0.45	78.27 +/- 0.11	79.10 +/- 0.09	78.83 +/- 0.24	79.47 +/- 0.20

 Table 8: Tabular data for the semi-supervised comparison shown in Fig. 1 for KNN ($k = 1$) classification accuracies on the Waveform dataset.

Labeled Data	5%	10%	20%	50%	70%	100%
ExpCLR (SS)	94.07 +/- 0.38	95.53 +/- 1.00	95.20 +/- 1.56	96.33 +/- 0.79	97.13 +/- 0.43	97.07 +/- 0.57
ExpCLR (S)	91.97 +/- 1.81	92.27 +/- 1.70	89.20 +/- 1.10	92.63 +/- 1.17	97.10 +/- 0.50	97.67 +/- 0.10
Cross-Entropy (S)	89.70 +/- 1.91	90.50 +/- 1.41	93.83 +/- 0.47	93.67 +/- 0.68	92.63 +/- 2.99	96.97 +/- 0.14
Kim et al. (2021) (SS)	78.17 +/- 4.75	78.67 +/- 6.27	78.40 +/- 3.58	89.10 +/- 3.42	90.33 +/- 3.12	90.97 +/- 2.26
SimCLR (SS)	91.43 +/- 2.35	93.90 +/- 1.53	93.80 +/- 1.03	93.90 +/- 1.26	94.70 +/- 1.21	95.90 +/- 1.35
SleepPriorCL (SS)	96.10 +/- 0.58	94.10 +/- 1.67	94.63 +/- 0.66	95.67 +/- 1.05	94.30 +/- 1.18	95.33 +/- 1.16

Table 9: Learning rates for each algorithm and dataset in unsupervised (U) and supervised (S) training mode, which we identified as best via a grid search. Note that some approaches reduce to a different loss during supervised training mode (see Sec. 4.3)

Dataset and lr	HAR (U)	HAR (S)	SleepEDF (U)	SleepEDF (S)	Waveform (U)	Waveform (S)
ExpCLR	3e-3	1e-3	5e-3	1e-2	5e-3	5e-3
Kim et al. (2021)	3e-3	1e-3	5e-3	1e-2	7e-3	5e-3
SimCLR	3e-3	5e-4	5e-3	1e-5	7e-3	5e-4
SleepPriorCL	1e-3	1e-3	7e-3	1e-3	1e-2	1e-3
Cross-Entropy	-	5e-4	-	1e-3	-	1e-3

$B_- \leq l_-$ and $B_+ \geq l_+$ for all encoders E_θ and linear maps M_ϕ which satisfy $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$. Next we set ϕ and θ to a value for which $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$, i.e. $M_\phi(E_\theta(x_i)) = f_i$ reproduces the expert features f_i on each datapoint x_i . Thus, by Def. 1 the best possible values for l_- and l_+ are

$$l_- = \min_{i,j} \frac{\|M_\phi(E_i - E_j)\|_2}{\|E_i - E_j\|_2} \quad \text{and} \quad l_+ = \max_{i,j} \frac{\|M_\phi(E_i - E_j)\|_2}{\|E_i - E_j\|_2}, \quad (6)$$

where i, j run over all pairs in $\{1, \dots, N\}$ with $E_i \neq E_j$. These values l_- and l_+ are strictly positive due to $\dim(\ker(M_\phi)) = 0$. Next, we use that the encoder is a neural network (NN) with L layers, so that the weights θ can be grouped as $\theta = (\theta_1, \dots, \theta_L)$, where θ_i are the weights of the i -th layer. Then define $\tilde{\theta} := (\theta_1, \dots, \theta_{L-1}, \frac{1}{c}\theta_L)$ and $\tilde{\phi} := c\phi$ for a constant $c \in \mathbb{R}_+$, where $\tilde{\phi}$ directly parameterizes the matrix $M_{\tilde{\phi}}$. Since $E_\theta \circ M_\phi = E_{\tilde{\theta}} \circ M_{\tilde{\phi}}$, it also holds that $\mathcal{L}_{mse}(M_{\tilde{\phi}} \circ E_{\tilde{\theta}}(X), F) = 0$ attains the global loss minimum, for any $c \in \mathbb{R}^+$. While this rescaling does not change the NN output as the two factors cancel, it changes l_- to $\tilde{l}_- = cl_-$ and l_+ to $\tilde{l}_+ = cl_+$. Thus, by choosing $c \in \mathbb{R}_+$ appropriately, one can attain any positive values for l_- or l_+ while still assuring that $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$. Therefore as claimed, it is not possible to provide any guarantees or bounds B_-, B_+ on l_-, l_+ solely from the condition $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$.

Next, let us consider the case $\dim(\ker(M_\phi)) > 0$. This can be either due to the specific choice of ϕ or hold generally when $e > d$ since one has $\dim(\ker(M_\phi)) \geq e - d > 0$. In this case the embedding is not unique for a given expert feature vector f_i since for any $\tilde{E} \in \ker(M_\phi)$ we have $M_\phi(E_i + \tilde{E}) = M_\phi(E_i) = f_i$. Now let us assume there exists a constant $l_- > 0$ such that for all embeddings which satisfy $M_\phi(E_i) = f_i \forall i$ (which is equivalent to $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$) it holds

that $l_- \|E_i - E_j\|_2 \leq \|f_i - f_j\|_2 \forall i, j$. Next, fix such an embedding E_θ and consider a fixed pair $k, l \in \{1, \dots, N\}$ with $f_k \neq f_l$. Then choose $\tilde{E} \in \ker(M_\phi)$ such that $\|\tilde{E}\|_2 \geq \frac{2}{l_-} \|f_k - f_l\|_2 + \|E_k - E_l\|_2$. Then a neural network embedding E' can be constructed that satisfies $E'(x_k) = E_k + \tilde{E}$ and $E'(x_i) = E_i \forall i \neq k$. Note that $M_\phi(E'(x_i)) = M_\phi(E_\theta(x_i)) \forall i \in \{1, \dots, N\}$, so that also $\mathcal{L}_{mse}(M_\phi \circ E'(X), F) = 0$. However, E' does *not* satisfy the aforementioned bilipschitz inequality for l_- since by the triangle inequality it holds

$$l_- \|E'(x_k) - E'(x_l)\|_2 = l_- \|E_k + \tilde{E} - E_l\|_2 \geq l_- \left(\|\tilde{E}\|_2 - \|E_k - E_l\|_2 \right) \geq 2\|f_k - f_l\|_2 > \|f_k - f_l\|_2,$$

therefore violating the inequality. Thus, solely from $\mathcal{L}_{mse}(M_\phi \circ E_\theta(X), F) = 0$ a non-zero value for l_- cannot be concluded, and thus it is not guaranteed that E_θ is a bilipschitz embedding. \square

C.2. Proof of Proposition 2

From $\mathcal{L}_{quad}(E(X), F) = 0$ it directly follows that $\forall i, j \in \{1, \dots, N\}$ we have $((1 - s_{ij})\Delta - D_{ij}) = 0$ and thus

$$\frac{\Delta}{\max_{k,l} \|f_k - f_l\|_2} \|f_i - f_j\|_2 = \|E_i - E_j\|_2.$$

Therefore, one can set $l_- = l_+ = \frac{\max_{k,l} \|f_k - f_l\|_2}{\Delta}$. \square

While the similarity measure used in Proposition 2 provides very good bounds on l_- and l_+ it is easy to see that other similarity measures like the one introduced in Sec. 3.6 also provide good bounds for l_+ and l_- .

C.3. Proof of Proposition 3

We start by proving (a) and then move on to (b). (a) can be easily shown by making use of the softmax limit

$$\tau \ln \sum_i \exp\left(\frac{a_i}{\tau}\right) \rightarrow \max_i a_i \quad \text{as } \tau \rightarrow 0,$$

for $a_i \in \mathbb{R}$. Using this one can directly see that

$$\tau \log \left[\frac{1}{N^2} \sum_{i,j} \exp(L_{ij}/\tau) \right] \rightarrow \max_{i,j} (L_{ij}).$$

For (b) and the limit $\tau \rightarrow \infty$, we start by using a series-expansion of $L_{ExpCLR}^\tau(E(X), F)$ which gives

$$L_{ExpCLR}^\tau(E(X), F) = \tau \log \left[\frac{1}{N^2} \sum_{i,j} \left(1 + \frac{L_{ij}}{\tau} + O\left(\frac{1}{\tau^2}\right) \right) \right] = \frac{1}{N^2} \sum_{i,j} L_{ij} + \mathcal{O}\left(\frac{1}{\tau}\right).$$

And thus for $\tau \rightarrow \infty$ we get $L_{ExpCLR}^\tau(E(X), F) \rightarrow \frac{1}{N^2} \sum_{i,j} L_{ij}$. \square

D. Statistical Bounds

The loss $L_{ExpCLR}^\tau(E(X), F)$ proposed by our work (see Eq. 4) mostly focuses on the pairs with largest loss value due to the hard negative mining (when τ is small). If we define the pair-Lipschitz constant to be $Z_{ij} := \frac{\|f_i - f_j\|_2}{\|E(x_i) - E(x_j)\|_2}$, the pairs with largest loss tend to be the ones whose pair-Lipschitz constant is furthest away from $\frac{\max_{k,l} \|f_k - f_l\|_2}{\Delta}$. Defining $l_{min} = \min_{k,l} Z_{kl}$ and $l_{max} = \max_{k,l} Z_{kl}$, we have $Z_{ij} \in [l_{min}, l_{max}] \forall i, j$; here i, j run over all indices in the respective training set. As explained in the methods section (Prop. 2), during training our loss tries to push l_{min} and l_{max} close to $\max_{kl} \|f_k - f_l\|/\Delta$.

While so far this only guarantees the pair-Lipschitz constants Z_{ij} from the training set to lie in the interval $[l_{min}, l_{max}]$, the question arises, whether we can guarantee that the Z_{ij} are also inside a certain interval for new unseen samples x_i, x_j .

We therefore split our dataset into two sets D_{train} and D_{val} of pairs of datapoints, where D_{train} consists of N_{train} pairs of i.i.d. samples and D_{val} of N_{val} pairs of i.i.d. samples. The pair-Lipschitz constant is thus a random variable, determined

by sampling two i.i.d. samples x_i, x_j from the underlying data distribution \mathbb{P} and then evaluating $Z_{ij} := \frac{\|f_i - f_j\|_2}{\|E(x_i) - E(x_j)\|_2}$, where f_i, f_j are the corresponding expert features of the input samples and $E(x_i), E(x_j)$ denotes the representations. The encoder E is updated by minimizing a loss function. We now present two approaches to obtain statistical bounds on the pair-Lipschitz constant for new unseen pairs of samples.

D.1. First Approach: Interval boundaries from validation set

- 1) Train an encoder E by minimizing the loss over D_{train} .
- 2) Calculate l_{min}^{val} and l_{max}^{val} as the minimum and maximum over the Z_{ij} on the validation set D_{val} .
- 3) Then we get via a PAC-bound that with prob. $(1 - \delta)$ w.r.t. repeated sampling of the validation set we have

$$P(Z_{test} \notin [l_{min}^{val}, l_{max}^{val}]) \leq \sqrt{\frac{8 \ln [2N_{val}(2N_{val} - 1) \frac{4}{\delta}]}{N_{val}}}, \quad (7)$$

where $Z_{test} := \frac{\|f - f'\|_2}{\|E(x) - E(x')\|_2}$ is the pair-Lipschitz constant obtained on two unseen (test) i.i.d. samples x and x' .

D.2. Second Approach: Interval boundaries from training set

- 1) Train an encoder E by minimizing the loss over D_{train} and calculate l_{min}^{train} and l_{max}^{train} as the minimum and maximum over the Z_{ij} on the training set D_{train} .
- 2) Calculate $P_{val} := \frac{1}{N_{val}} \sum_{(i,j) \in D_{val}} 1_{Z_{ij} \notin [l_{min}^{train}, l_{max}^{train}]}$ on D_{val} .
- 3) Then again via a PAC-bound we get that with prob. $(1 - \delta)$ w.r.t. repeated sampling of the validation set

$$P(Z_{test} \notin [l_{min}^{train}, l_{max}^{train}]) \leq P_{val} + \sqrt{\frac{\ln [\frac{2}{\delta}]}{N_{val}}}, \quad (8)$$

where again $Z_{test} := \frac{\|f - f'\|_2}{\|E(x) - E(x')\|_2}$ is the pair-Lipschitz constant obtained on two unseen (test) i.i.d. samples x and x' .

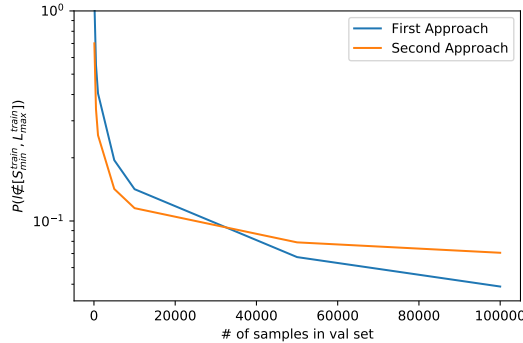


Figure 8. Empirical comparison of the two bounds from Apps. D.1 resp. D.2: The figure shows the upper bounds on the probability that the pair-Lipschitz constant lies outside $[l_{max}, l_{min}]$ for a randomly drawn test pair, over the number of pairs in the validation set N_{val} . The empirical evaluation shows that for a large number of validation pairs the first approach leads to better guarantees, while for smaller number of pairs the second approach beats the first approach.

D.3. Empirical Evaluation of the Bounds

We now empirically evaluate which of the two previous bounds gives better guarantees. Therefore, we assume to have a discrepancy between training and validation bounds which leads to 5% of the validation pairs to have sampled Z_{ij} outside of the training bounds $[l_{min}^{train}, l_{max}^{train}]$ and also that we have some overfitting which leads to a smaller interval on the training set. Thus, for illustration purposes, we assume $P_{val} \approx 0.05$, $l_{min}^{val} \leq l_{min}^{train}$ and $l_{max}^{val} \geq l_{max}^{train}$.

We evaluate both bounds for different numbers of pairs in the validation set N_{val} and calculate our bounds on the probability of a sampled Z_{test} , generated by a pair of i.i.d. samples, to lie outside the window $[l_{max}, l_{min}]$. (Here we consider the case where we are mostly interested in the probability bounds, not so much in the exact interval boundaries.) The comparison over the number of validation samples is shown in Fig. 8.

If one only requires an upper *or* a lower bound on the pair-Lipschitz constant, the bound can be improved since as growth function reduces to $G(n) = n$. This then improves the bound from App. D.1 slightly: In this case with probability at least $(1 - \delta)$ w.r.t. repeated sampling of the validation set we have

$$P(Z_{test} \geq l_{max}^{val}) \leq \sqrt{\frac{8 \ln \left[\frac{8N_{val}}{\delta} \right]}{N_{val}}}.$$

E. Comparison of Loss Function Variants and their Derivatives

Below one can find the equations for ExpCLR (NHNM) and the pair loss with their respective derivatives. Fig. 9 additionally visualizes them over a range of values for the distance metric D_{ij} . One can see that ExpCLR has the same minimum as the pair loss, but possesses continuous derivatives w.r.t. D_{ij} .

$$\mathcal{L}_{ij,ExpCLR(NHNM)} = ((1 - s_{ij})\Delta - D_{ij})^2$$

$$\frac{\partial \mathcal{L}_{ij,ExpCLR(NHNM)}}{\partial D_{ij}} = -2((1 - s_{ij})\Delta - D_{ij})$$

$$\mathcal{L}_{ij,pair} = s_{ij}D_{ij}^2 + \max(0, (1 - s_{ij})^2\Delta^2 - D_{ij}^2)$$

$$\frac{\partial \mathcal{L}_{ij,pair}}{\partial D_{ij}} = \begin{cases} 2s_{ij}D_{ij} & D_{ij} \geq \Delta(1 - s_{ij}) \\ -2D_{ij}(1 - s_{ij}) & \text{else} \end{cases}$$

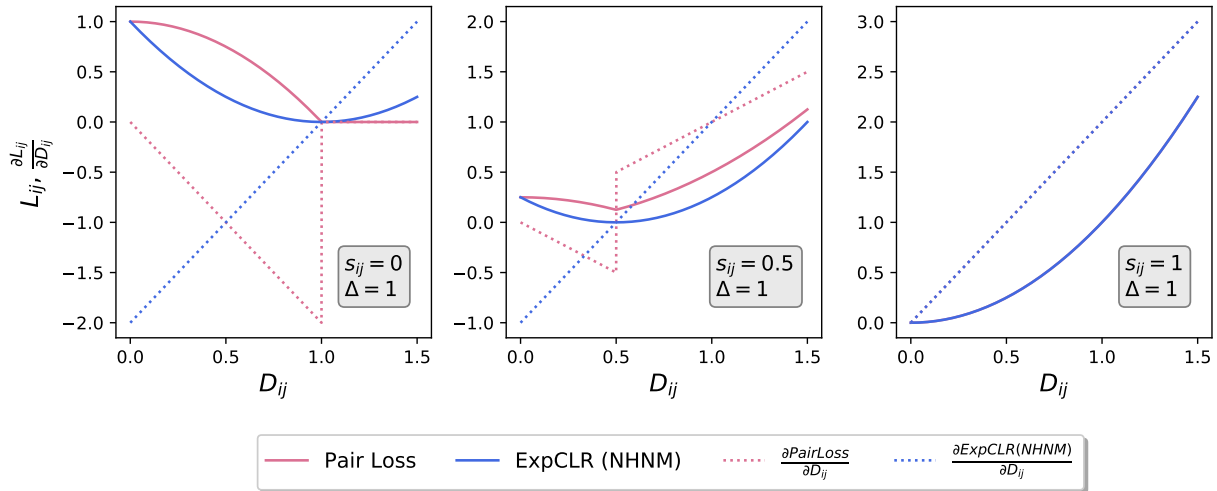


Figure 9. Comparison of ExpCLR (NHNM) and the Pair-Loss with their associated Derivatives: This figure shows the Loss function of ExpCLR (NHNM=no hard-negative mining) from Eq. 3 and the pair-loss from Eq. 1 and their derivatives. On the y-axis one can see values for the loss function resp. derivative of the loss function and on the x-axis certain values for the distance metric D_{ij} are given. From the left to the right panel different scalar values of the similarity measure s_{ij} are shown. Note that in right plot both losses reduce to the same as $s_{ij} = 1$.

F. Gradient-Level Proof of Hard-Negative Mining

In this section we demonstrate on gradient-level how our hard-negative mining scheme introduced in Sec. 3.5 functions. This complements the intuitions provided by Prop. 3. For this, we take the gradient of Eq. 4:

$$\begin{aligned} \frac{\partial}{\partial L_{nm}} \mathcal{L}_{ExpCLR}^{\tau}(E(X), F) &= \frac{\partial}{\partial L_{nm}} \tau \log \left[\sum_{i,j=1}^N \frac{\exp\left(\frac{L_{ij}}{\tau}\right)}{N^2} \right] \\ &= \exp\left(\frac{L_{nm}}{\tau}\right) \left[\sum_{i,j=1}^N \exp\left(\frac{L_{ij}}{\tau}\right) \right]^{-1}. \end{aligned}$$

Therefore, the gradient is directly proportional to $\exp\left(\frac{L_{nm}}{\tau}\right)$, adding an exponential (softmax) scaling and thus increasing the gradient contributions of those pairs (n, m) with larger loss components L_{nm} .