
SOM-CPC: Unsupervised Contrastive Learning with Self-Organizing Maps for Structured Representations of High-Rate Time Series

Iris A.M. Huijben^{1,2} Arthur A. Nijdam¹ Sebastiaan Overeem^{1,3} Merel M. van Gilst^{1,3} Ruud J.G. van Sloun¹

Abstract

Continuous monitoring with an ever-increasing number of sensors has become ubiquitous across many application domains. However, acquired time series are typically high-dimensional and difficult to interpret. Expressive deep learning (DL) models have gained popularity for dimensionality reduction, but the resulting latent space often remains difficult to interpret. In this work we propose SOM-CPC, a model that visualizes data in an organized 2D manifold, while preserving higher-dimensional information. We address a largely unexplored and challenging set of scenarios comprising high-rate time series, and show on both synthetic and real-life data (physiological data and audio recordings) that SOM-CPC outperforms strong baselines like DL-based feature extraction, followed by conventional dimensionality reduction techniques, and models that jointly optimize a DL model and a Self-Organizing Map (SOM). SOM-CPC has great potential to acquire a better understanding of latent patterns in high-rate data streams.

1. Introduction

The improvement and abundance of sensor technology has led to large amounts of high-dimensional, information-rich continuous data streams. However, gaining actionable insights from these data is challenging due to their low interpretability. The main objective of this study is to develop an algorithm for acquiring a structured and interpretable representation of high-rate time series. We define high-rate time series as data streams that are sampled

¹Department of Electrical Engineering, Eindhoven University of Technology, The Netherlands ²Onera Health, Eindhoven, The Netherlands ³Sleep Medicine Center Kempenhaeghe, Heeze, The Netherlands. Correspondence to: Iris Huijben <i.a.m.huijben@tue.nl>.

at their Nyquist rate (often hundreds of samples per second), which are distinct from, e.g., (medical) tabular data, which contain sparsely sampled information (e.g. once per hour). We specify an interpretable representation as one that has the ability to be informative and to facilitate exploration of the underlying structure over time (Lipton, 2018).

According to the manifold hypothesis, high-dimensional real-world data lies on a low-dimensional manifold, comprising disentangled latent factors of variation. Dimensionality reduction techniques like Principle Component Analysis (PCA) have conventionally been used to reveal this manifold. However, acquiring an interpretable representation with PCA typically requires omitting many principle components, which may discard important information that can not be linearly projected on the few remaining components. Moreover, additional clustering methods like K-means clustering, are typically needed to find structures or clusters in the the resulting low-dimensional projections.

A Self-Organizing Map (SOM) (Kohonen, 1990) is an extension of K-means clustering that can be used without the need to first project to the low-dimensional manifold. It creates a low-dimensional interpretable visualization, while still representing the data in multiple dimensions. However, like PCA and K-means, SOMs typically act on features, which need to be selected heuristically and may, therefore, strongly depend on the use case and/or data modality.

Deep learning (DL) models have become popular general data-driven feature extractors, and the sub-area of unsupervised representation learning, is concerned with models that learn to represent data with a set of features, learned without the bias of human annotations. To enhance interpretability, latent space representations of such models are often visualized using a t-distributed stochastic neighbor embedding (t-SNE) (Hinton & Roweis, 2002). Albeit its frequent use, t-SNE does not allow a direct deployment on unseen data as it does not learn a reusable mapping between the multi-dimensional and the low-dimensional space.

DL models have, therefore, also been combined with (joint) clustering objectives in the latent space. Resulting deep-clustering methods are, however, often designed

for static data - and if tested on time series - assign a cluster to a full time series rather than to subsequent windows (Xie et al., 2016; Yang et al., 2017), preventing exploration of temporally-changing patterns. Also, some require an additional step to create a visually interpretable representation (Madiraju, 2018), or make use of labels to improve training (Lee & Schaar, 2020). These labels are by definition unavailable for (yet) hidden patterns.

A sub-area of deep-clustering methods has focused on regularizing autoencoder training with a SOM objective in the latent space (Ferles et al., 2018; Pesteie et al., 2018; Fortuin et al., 2019; Forest et al., 2019; Manduchi et al., 2021; Forest et al., 2021). However, similar to Mrabah et al. (2020), we hypothesize that the reconstruction objective of an autoencoder may hamper the clustering or structured representation learning objective: while within-cluster similarities should remain preserved for latent clustering (e.g. preserving only the underlying state of a physiological signal), reconstruction demands a preservation of within-cluster differences as well (e.g. the high-frequency noise). Moreover, in the context of time series representation learning, other self-supervised models - that exploit the temporal dimension - might be more suitable.

Contrastive self-supervised learning approaches have quickly become popular thanks to their superior representation learning performance in many domains (see (Le-Khac et al., 2020) for a review). While many of these models rely on data augmentations during training in order to construct pairs of similar data points, Contrastive Predictive Coding (CPC) (Oord et al., 2019) leverages the temporal dimension for this purpose, making it a natural choice for self-supervised representation learning of time series. In CPC, the temporal dimension not only serves as a pretext task, but simultaneously enforces latent smoothness over time. The contributions of this work are as follows:

- We formulate a general framework that captures all previous deep-SOM literature, and inherit it with SOM-CPC, a model for learning structured and interpretable 2D representations of (high-rate) time series by encoding subsequent data windows to a topologically ordered set of quantization vectors.
- We show that SOM-CPC preserves more information in its 2D representation and exhibits more temporal smoothness in 2D than CPC followed by PCA, a linear classifier or K-means, or directly encoding CPC’s latent space to two dimensions.
- We show that SOM-CPC quantitatively and qualitatively outperforms autoencoder-based deep-SOM models in terms of both clustering and topological ordering, while requiring less auxiliary loss functions and associated hyperparameter tuning.

2. Preliminaries

2.1. Kohonen Self-Organizing Maps

Kohonen’s Self-Organizing Map (SOM) (Kohonen, 1990) is an algorithm to find a visually interpretable topological data representation. It has been found useful to reveal intricate patterns and structure in a plethora of applications (Yin, 2008). The algorithm’s output, the low-dimensional visualization, is often referred to as a SOM as well.

We define a data point $z \in \mathcal{Z}$, and its quantized counterpart $q_{\Phi}(z) \in \Phi$, with $\Phi = [\phi_1; \dots; \phi_k]$ a trainable quantization codebook containing k vectors or prototypes $\phi_i \in \mathbb{R}^F$, $1 \leq i \leq k$. Vector $q_{\Phi}(z) := \Phi[\operatorname{argmin}_i (\|\phi_i, z\|_2^2)]$ is thus the ‘winning quantization vector’. Note that Φ is updated every training iteration n , which we omit for readability.

The codebook vectors are placed on a pre-defined 2D grid of nodes by assigning an xy-coordinate to each vector at initialization. This creates a 2D representation, while each data point z lives in \mathbb{R}^F , with $F \gg 2$. This is conceptually different than how PCA achieves dimensionality reduction to 2D, where all information in the 3rd and higher principle components is strictly omitted. During SOM training, each ϕ_i is updated as follows (Kohonen, 1990):

$$\phi_i^{(n+1)} = \phi_i^{(n)} + \eta^{(n)} \mathcal{S}_i(q_{\Phi}(z))(z - \phi_i^{(n)}), \quad (1)$$

where $\eta^{(n)}$ is an decreasing learning rate. Topological neighborhood structure is promoted via a neighbourhood kernel \mathcal{S} that weighs nodes inversely proportional to their distance with the winning node. A Gaussian kernel is often used, which weighs node i according to:

$$\mathcal{S}_i(q_{\Phi}(z)) = \exp\left(-\frac{d_i^{(n)}}{2(\sigma^{(n)})^2}\right), \quad \text{with} \quad (2)$$

$$d_i^{(n)} = \|\mathcal{P}[q_{\Phi}(z)], \mathcal{P}[\phi_i^{(n)}]\|_2^2 \quad \text{and} \quad (3)$$

$$\sigma^{(n)} = \sigma^{(0)} \exp(-n/\lambda), \quad (4)$$

where \mathcal{P} projects a codebook vector to its corresponding coordinate on the grid, $\sigma^{(0)}$ denotes the initial standard deviation, and λ a decay factor. The dependence of \mathcal{S}_i on d_i , implies a weighing of 1 for the winning node (i.e. distance equals zero), and lower than 1 for neighbour nodes. Other neighbourhood structures have been proposed as well, for example using the four closest neighbours on the grid, resulting in a kernel with a *plus*-shape (Fortuin et al., 2019).

2.2. Deep-SOM models

Deep-SOM research has focused on combining autoencoders (Ferles et al., 2018; Pesteie et al., 2018; Fortuin et al., 2019; Forest et al., 2019; Manduchi et al., 2021; Forest et al., 2021) with a SOM. These models can broadly be summarized as a vector-quantized variational

autoencoder (VQ-VAE) (van den Oord et al., 2017), with a topological organization of the vectors in the quantization codebook: the SOM. The models are trained end-to-end using error backpropagation of both a reconstruction *task* loss $\mathcal{L}_{\text{task}}$ and a loss $\mathcal{L}_{\text{topo}}$ that encourages *topological* ordering in the SOM. In general, a deep-SOM training objective takes the following form:

$$\mathcal{L}_{\text{deep-SOM}} = \mathcal{L}_{\text{task}} + \alpha \mathcal{L}_{\text{topo}}, \quad \text{with} \quad (5)$$

$$\mathcal{L}_{\text{topo}}(\mathbf{z}^{(n)}) = \mathbb{E}_{\mathcal{Z}} \left[\sum_{i=1}^k \mathcal{S}_i(q_{\Phi}(\mathbf{z}^{(n)})) \|\mathbf{z}^{(n)} - \phi_i^{(n)}\|_2^2 \right]. \quad (6)$$

Hyperparameter α controls the trade-off. The topological loss replaces the original update rule of the SOM algorithm from Equation (1), and is in practice computed on mini-batches of data to approximate the expectation. The features in \mathcal{Z} are jointly optimized via the parameters of the encoder, and thus also depend on n now. To prevent clutter we will, however, omit the (n)-superscript in the following.

Fortuin et al. (2019) propose SOM-VAE. As opposed to VQ-VAE, SOM-VAE has two decoders, as it also decodes the continuous latents. The topological loss was split in a *commitment* loss (committing the winning codebook vector to \mathbf{z} and *vice versa*) and a *SOM* loss (pulling the codebook vectors of the neighbours to \mathbf{z}): $\mathcal{L}_{\text{topo}} = \mathcal{L}_{\text{commitment}} + \frac{\beta}{\alpha} \mathcal{L}_{\text{SOM}}$. Formally:

$$\mathcal{L}_{\text{commitment}} = \mathbb{E}_{\mathcal{Z}} \left[\|\mathbf{z} - q_{\Phi}(\mathbf{z})\|_2^2 \right], \quad \text{and} \quad (7)$$

$$\mathcal{L}_{\text{SOM}} = \mathbb{E}_{\mathcal{Z}} \left[\sum_{\substack{i=1; \\ \phi_i \neq q_{\Phi}(\mathbf{z})}}^k \mathcal{S}_i(q_{\Phi}(\mathbf{z})) \|\text{sg}[\mathbf{z}] - \phi_i\|_2^2 \right], \quad (8)$$

with $\text{sg}[\cdot]$ a gradient blocker, making the encoder parameters independent of the quantization error of the neighbour nodes. The authors chose a plus-shaped neighbourhood kernel \mathcal{S} . Note that for $0 < \beta/\alpha < 1$, this plus-kernel is a coarse approximation of the Gaussian kernel. The sum of the reconstruction losses $\mathcal{L}_{\text{recon,cont}}$ and $\mathcal{L}_{\text{recon,disc}}$ from the continuous and discrete decoder, respectively, yield the total task loss $\mathcal{L}_{\text{task}}$, which is combined with the topological loss to create the training objective of the SOM-VAE model.

SOM-VAE-prob (Fortuin et al., 2019) and (T)-DPSOM (Manduchi et al., 2021) are extensions of SOM-VAE. SOM-VAE-prob enforces smoothness over time by adding a transition loss (multiplied by γ) to optimize a first-order Markov model to learn the node transition probabilities, and a smoothness loss (multiplied by τ in their work, we will use ζ) to minimize the quantization error of highly probable transitions. DPSOM is a probabilistic model, based on a VAE with a non-degenerate approximate posterior (Kingma

& Welling, 2013) with soft cluster assignment and a cluster assignment hardening (CAH) loss (Xie et al., 2016). T-DPSOM additionally incorporates a temporal smoothness loss, and an LSTM, which aims to predict the future latent space. Note that the measures for additional temporal smoothness in SOM-VAE-prob and T-DPSOM impose additional loss-weighting hyperparameters that need tuning, possibly for each data modality separately. This smoothness is already naturally embedded in the task objective of our model (see Section 3.2). The probabilistic additions in the (T)-DPSOM model, with respect to SOM-VAE, are orthogonal to the developments in this work.

Forest et al. (2021) propose Deep embedded SOM (DESOM). Compared to SOM-VAE, the decoder on the discrete space is omitted (therewith also $\mathcal{L}_{\text{recon,disc}}$), gradients from \mathcal{L}_{SOM} to the encoder are not being blocked (i.e. $\text{sg}[\cdot]$ is removed from Equation (8)), and the topological loss $\mathcal{L}_{\text{topo}}$ is given by Equation (2), i.e. with a Gaussian neighbourhood function with decaying variance. In a short work, Forest et al. (2019) speculate about adding an LSTM in the latent space to train a SOM on sequential data, and refer to this model as LSTM-DESOM. Figure 1a-b visualizes the SOM-VAE and DESOM architecture.

3. SOM-CPC

3.1. Motivation

Discovering new patterns in high-dimensional data using SOMs requires features that are both suitable for SOM organization *and* accurately reflect the data. A desirable property of a feature extractor is that it inverts the data-generating process, inferring a latent low-dimensional state that carries useful information. This data-generating process is in general unknown, implicit and highly non-linear. Feature learning can thus be formulated as a non-linear independent component analysis problem, which was proven to be non-identifiable (Hyvärinen & Pajunen, 1999). However, recent advances showed that the problem becomes identifiable under the assumed presence of an auxiliary variable (Hyvärinen et al., 2018). Such an auxiliary variable (e.g. a temporal component) is not present in plain autoencoders, but the contrastive learning paradigm was shown to conform to this assumption (Hyvärinen et al., 2018; Zimmermann et al., 2021).

In line with that, Oord et al. (2019) showed that minimizing the InfoNCE loss in the contrastive learning framework CPC maximizes the mutual information between a context vector that summarizes past and current information, and a latent space from a future data window. CPC thus results in features that encode a slowly-changing underlying state. Encoding (high-frequency) information that is less-shared across windows is thus not needed to minimize InfoNCE,

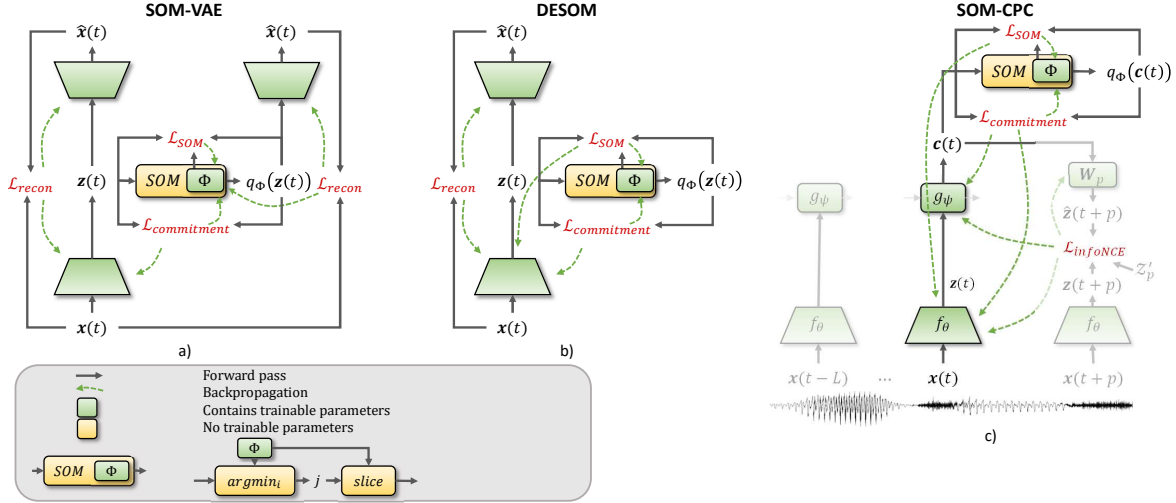


Figure 1: Architectures of different deep-SOM models including the gradient paths in green. a) SOM-VAE (Fortuin et al., 2019) b) DESOM (Forest et al., 2021) c) SOM-CPC (ours). The two decoders in the SOM-VAE model are independent and have their own trainable parameters, while the visualized encoders in the SOM-CPC model are all the same (i.e. parameters are shared). The g_ψ block in the SOM-CPC model indicates an autoregressive component (e.g. a GRU), and \mathcal{Z}'_p refers to a set of drawn negative embeddings.

while such within-cluster/state differences typically do need to be encoded for reconstruction. This is hypothesized to make autoencoders less suitable to do feature extraction for clustering purposes (Mrabah et al., 2020).

We thus propose to leverage CPC as a feature extractor, and jointly optimize it with a SOM, which results in SOM-CPC: a representation learning model that learns to map windows of time series data to a structured 2D grid for the purpose of pattern discovery.

3.2. Algorithmic Details

We introduce $\mathbf{X} = [\dots, \mathbf{x}(t), \mathbf{x}(t+1), \dots]$, a sequence of non-overlapping data windows $\mathbf{x}(t) \in \mathbb{R}^{ch \times T}$, with ch the number of channels, T the number of samples in a window, and t the index of the window. For brevity we omit this window index when possible. Following Oord et al. (2019), an encoder, parameterized by θ , maps each window \mathbf{x} to a latent representation $\mathbf{z} = f_\theta(\mathbf{x}) \in \mathbb{R}^F$, with F the number of features. \mathcal{Z} includes the embeddings of all windows in \mathbf{X} . The sets containing all available data streams \mathbf{X} and corresponding embeddings \mathcal{Z} are denoted with \mathcal{X} and \mathcal{Z} , respectively. A causal auto-regressive (AR) module g_ψ parameterized by ψ , e.g. a gated-recurrent unit (GRU), subsequently aggregates the current and L previous embeddings into a context vector $\mathbf{c}(t) \in \mathbb{R}^F$. From this context, predictions are made for P future (or ‘positive’) embeddings $\mathbf{z}(t+p)$, with $p \in \{1, \dots, P\}$. The task objective in SOM-CPC maximizes the match (expressed as a dot-product) between these predictions and the positive embeddings, as compared to this match for N ‘negative’ embeddings. These negatives may be sampled across the dataset (i.e. from \mathcal{Z}), or within the same time-series (i.e.

from \mathcal{Z}). The task objective, being the InfoNCE loss (Oord et al., 2019), is defined as:

$$\mathcal{L}_{\text{task}} := \mathcal{L}_{\text{InfoNCE}} = \frac{1}{P} \sum_{p=1}^P \mathcal{L}_p, \quad \text{with } \mathcal{L}_p = \quad (9)$$

$$- \mathbb{E}_{\mathcal{X}} \left[\log \frac{\exp(\mathbf{z}(t+p) \mathbf{W}_p \mathbf{c}(t))}{\sum_{\mathbf{z}' \in \mathcal{Z}'_p \cup \{\mathbf{z}(t+p)\}} \exp(\mathbf{z}' \mathbf{W}_p \mathbf{c}(t))} \right], \quad (10)$$

with $\mathcal{Z}'_p \subset \mathcal{Z}$ the embeddings of drawn negative samples ($|\mathcal{Z}'_p| = N$), and $\mathbf{W}_p \in \mathbb{R}^{F \times F}$ a trainable linear predictor between the current context vector and the p^{th} future embedding.

The context vector is not only used to predict future embeddings, it is also the input to the SOM module that selects the winning node. We use a 2D (square) grid for the SOM nodes to acquire visual interpretability. The SOM is optimized using the topological loss $\mathcal{L}_{\text{topo}}$, as defined in Equation (6), with a Gaussian neighbourhood kernel \mathcal{S} (see Equation (2)). Depending on the use case, it was found to not always be necessary, or even beneficial (due to higher risk of overfitting), to use an AR module g_ψ to aggregate causal context into the current embedding. If the AR module is not used, the future predictions are made directly from the current (continuous) latent space $\mathbf{z}(t)$ instead of $\mathbf{c}(t)$. Likewise, $\mathbf{z}(t)$ rather than $\mathbf{c}(t)$ is being quantized by the SOM module. Depending on the presence of this AR module, both $\mathcal{L}_{\text{topo}}$ and $\mathcal{L}_{\text{task}}$ are thus computed on either $\mathbf{z}(t)$ or $\mathbf{c}(t)$.

All model elements are jointly optimized, with the training objective: $\mathcal{L}_{\text{SOM-CPC}} = \mathcal{L}_{\text{task}} + \alpha \mathcal{L}_{\text{topo}}$, which adheres to the general objective of a deep-SOM model as formulated

in Equation (5). Figure 1c visualizes SOM-CPC, and its gradient paths in green. The initial standard deviation of the Gaussian neighbourhood kernel was set to $\sigma^{(0)} = \sqrt{k}/2$, ensuring that the full grid is part of the neighbourhood at the start of training. Pseudo-code can be found in Appendix A.1 (Algorithm 1) and code is published on <https://github.com/IamHuijben/SOM-CPC.git>.

3.3. Performance Evaluation

Forest et al. (2020) provide a taxonomy of SOM metrics that distinguishes *external vs internal* and *topological vs clustering* metrics. External metrics are related to labels (which are not used during unsupervised training), while internal metrics do not depend on such information. Topological metrics assess the topological ordering (i.e. neighbourhood relations) of the SOM, while clustering metrics are more related to, for example, pureness of nodes.

To evaluate clustering performance, linked to external labels, we leverage purity and the normalized mutual information (NMI). The latter corrects for a high number of clusters (i.e. nodes), which could easily lead to high pureness, but leaves the NMI more conservative.

Even though scoring high on external metrics is not the main goal of a representation learning model like SOM-CPC, we do report them as it provides an indication of how well information was preserved. To compute regression/classification performance, we first ‘color’ (or label) each node with the most occurring (for discrete labels) or median (for continuous labels) label from the training set. The test set predictions are then converted from node indices to label predictions by using these colorings. Regression performance is expressed as the average squared regression error with the target: SE_{target} . Classification performance is reported with Cohen’s kappa (Cohen, 1960), a commonly used metric that corrects for correctness by chance.

Topographic performance is measured using the (internal) Topographic Error (TE) (Kiviluoto, 1996), which reports the fraction of windows (between 0 and 1) for which the winning and second-best winning node are not neighbours in the SOM (lower is better). Finally, to measure whether a time series conveys a smooth trajectory through SOM space, we measure the average Euclidean distance (denoted $\ell_{2,\text{smooth}}$) between all subsequent windows in each time series. The lower this value, the less frequently large jumps in the 2D map occur. Note that in extreme cases where many windows collapsed to the same node, both the TE and the average $\ell_{2,\text{smooth}}$ metric are artificially pushed down. We can thus only interpret these metrics in conjunction with earlier-mentioned clustering and classification metrics. Mathematical definitions of all metrics are provided in Appendix A.1.

4. Experiments

SOM-CPC is compared to several 2D representation learning methods. First, deep-SOM models with a reconstruction task loss (i.e. SOM-VAE(-prob), and (GRU-)DESOM). Second, vanilla CPC with a high-dimensional latent space ($F \gg 2$) (Oord et al., 2019), followed by PCA for dimensionality reduction to 2D. Third, CPC with a 2D latent space ($F = 2$). For the CPC-based models, linear and non-linear read-out is, respectively, tested using a linear neural classifier, and K-means clustering with the same number of clusters as the number of nodes used in SOM-CPC. High-dimensional vanilla CPC ($F \gg 2$) without additional dimensionality reduction is, moreover, tested as it sets an upper bound for the amount of information that can be preserved given the encoder architecture, while not providing an interpretable 2D representation. The same encoder architecture is used for all models that are compared in a single application domain, and all models are run with the same seed for randomization. We benchmarked our implementation of SOM-VAE against reported results by Fortuin et al. (2019) (see Appendix A.2). Details on model architectures and training settings for the different applications can be found in Appendix A.3.1, A.4.2, and A.5.1.

4.1. Synthetic Data

Data generation: A synthetic dataset was created, consisting of sinusoids with an initial frequency sampled from a uniform distribution between 20 and 40 Hz. The frequency of the signals was altered over time according to a random walk process with a step size of 0.1 Hz. As such, at each time step (i.e. sample), the signal’s frequency either increased with 0.1 Hz (with probability $p_{\text{up}} = 0.1$), decreased with 0.1 Hz ($p_{\text{down}} = 0.1$), or remained constant ($p_{\text{c}} = 0.8$). In case the random walk crossed either 1 or 60 Hz, the probabilities were (temporarily) altered to $[p_{\text{up}}, p_{\text{c}}, p_{\text{down}}] = [0.5, 0.5, 0]$ or $[p_{\text{up}}, p_{\text{c}}, p_{\text{down}}] = [0.0, 0.5, 0.5]$, respectively. All series were finally corrupted with an additive white Gaussian noise vector $\epsilon \sim \mathcal{N}(0, 0.01)$. Formally, each generated signal took the form: $x[n] = \sin\left(2\pi \frac{f[n-1] + \Delta f}{f_s} n\right) + \epsilon$, with $f[n = 0] \sim U[20, 40]$, $\Delta f \sim \text{Categorical}([p_{\text{up}}, p_{\text{c}}, p_{\text{down}}])$, and $f_s = 128$ Hz the sampling frequency. A total of 200 of such time series, each of 5 minutes, were generated, and labels were defined per 1-second window by taking the median frequency. The set was randomly divided into a training ($n = 100$), validation ($n = 50$), and test split ($n = 50$).

Comparison to deep-SOM baselines: Table 1 shows that SOM-CPC outperforms all deep-SOM baselines on all metrics. Moreover, Figure 6 in Appendix A.3.2 reveals that the latent space of SOM-CPC, when projected on its two principle components, shows a better organization of the

Table 1: Mean and one std. dev. across all test set series of synthetic data. Results for varying values of α , γ , and ζ can be found in Table 3, Appendix A.3.2. SOM-CPC clearly outperforms all baselines. Models below the dashed line are ablations. Regression plots and SOMs of models with a * are visualized in Figure 2a-c.

Model	α	\mathcal{S}	$\mathcal{L}_{\text{SOM sg}}[\cdot]$	$\text{SE}_{\text{target}}$	$\ell_{2,\text{smooth}}$	TE
CPC + linear classifier	-	-	-	2.62± 2.37	-	-
CPC + K-means	-	-	-	1.09± .62	-	-
CPC ($F = 2$) + linear classifier	-	-	-	25.01±42.94	-	-
CPC ($F = 2$) + K-means	-	-	-	.76± 1.31	-	-
CPC + PCA + linear classifier	-	-	-	42.81±58.12	-	-
CPC + PCA + K-means	-	-	-	4.42± 9.01	-	-
* SOM-VAE	.1	Plus	✓	8.02± 4.58	2.41±.68	.28±.07
SOM-VAE	1e-3	Gaussian	✓	11.60±25.43	1.92±.34	.06±.02
SOM-VAE-prob	.1	Plus	✓	20.10±48.80	3.15±.73	.63±.05
* DESOM	.1	Gaussian	✗	10.77±10.85	2.20±.44	.06±.02
* SOM-CPC (ours)	1e-4	Gaussian	✗	.72± 1.08	1.37±.37	.02±.01
Ablations						
SOM-CPC	1e-2	Gaussian	✓	.47± .48	.99±.24	.07±.04
SOM-CPC	1e-4	Plus	✗	1.71± 1.15	2.46±.51	.30±.06
SOM-CPC	1e-2	Plus	✓	1.16± .61	1.85±.26	.12±.04
CPC + SOM (disjoint)	-	Gaussian	-	.84± 1.14	1.47±.50	.03±.01

signal frequency than the projections of the SOM-VAE and DESOM models. Figure 2a-c display the resulting SOMs (colored with the median test set labels) for the same three models. Uncolored nodes in the SOM were not assigned in the test set. Interestingly, although the SOM for the DESOM and SOM-CPC model look similar, the $\text{SE}_{\text{target}}$ is higher for the DESOM model, which can also be seen from the regression plots below the SOMs.

The addition of two temporal losses in the SOM-VAE-prob model, as compared to SOM-VAE, did deteriorate the given metrics, even though a range of values for multipliers α , γ and ζ was tested (see Table 3 in Appendix A.3.2 for the full sweep). The deterioration of the results can be explained by the difficulty of finding the correct scaling factors for these additional losses. Note that the SOM-CPC model automatically incorporates smoothness over time thanks to the nature of the CPC task loss, therewith preventing additional hyperparameter tuning.

Additionally, we study the optimization behaviour of SOM-VAE, DESOM, and SOM-CPC by plotting the progression of the task versus the topological loss during training (see Figure 2d). To make a fair comparison, we plot the SOM-VAE models that are trained with a Gaussian neighbourhood kernel. Different curves in the graphs indicate runs with varying values for α , and the line color’s gradient denotes the training iteration. The SOM-CPC graphs include the models run with and without gradient detachment of \mathcal{L}_{SOM} to the encoder. It can be seen that both losses jointly minimize in SOM-CPC training, while there is a counteracting effect visible for SOM-VAE, and a high influence of the value of α for DESOM training.

Comparison to other baselines: CPC (with $F = 2$), and CPC followed by PCA, resulted in a much higher regression error $\text{SE}_{\text{target}}$ than SOM-CPC when using linear read-out (see Table 1). Non-linear K-means clustering improved performance for both cases, but only for CPC with $F = 2$, performance nearly reached SOM-CPC performance. Later we will see that optimizing CPC with $F = 2$ can hamper optimization for more intricate data spaces (see Section 4.3). Interestingly, regression performance of SOM-CPC was found to be even slightly better than that of the vanilla multi-dimensional CPC model (with $F = 128$), both for linear classification and K-means. This could be explained by the additional regularization that the SOM provides in SOM-CPC training.

Ablations: We perform several ablation experiments on SOM-CPC, which are reported below the dashed line in Table 1. Blocking the gradients of the neighbour nodes with respect to the encoder during training ($\mathcal{L}_{\text{SOM sg}}[\cdot]$ column) slightly improved the regression error and temporal smoothness, but decreased the topographic error. Looking at models run with various values for α (reported in Table 3, Appendix A.3.2), the effect of gradient blocking can be considered small and ambiguous when considering different metrics. Disjoint training (CPC + SOM) resulted in a less smooth trajectory over time through the 2D SOM space, seen from the higher $\ell_{2,\text{smooth}}$. Using a plus neighbourhood kernel instead of a Gaussian kernel decreased performance on all three metrics. The increase in TE, is well explainable by the fact that a plus kernel takes into account fewer neighbours (at least at the start of training) and therefore has more difficulty to find a good topological mapping. Figure 3 shows the development of the SOM node spread (projected

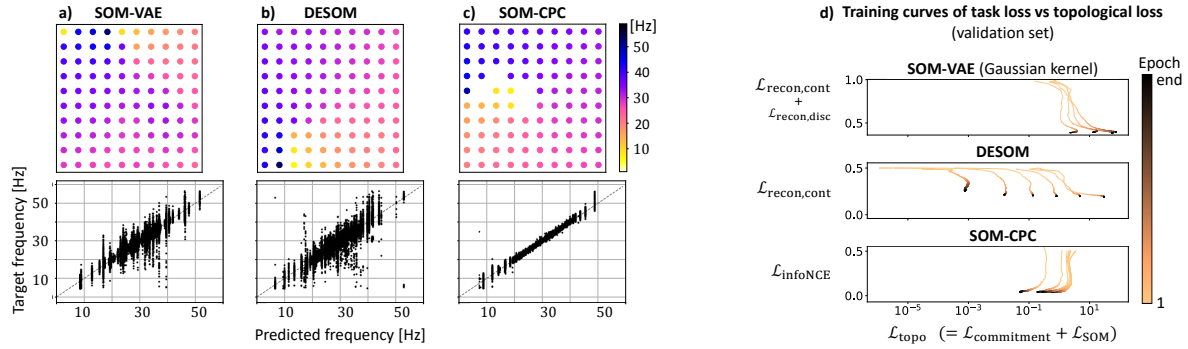


Figure 2: SOMs and regression plots for SOM-VAE (a), DESOM (b) and SOM-CPC (c). Both DESOM and SOM-CPC show a gradual change of frequency over the grid, but the regression error SE_{target} is lower for SOM-CPC, which can also be seen from the regression plot, where the predicted window frequencies are plotted against the target frequencies (i.e. training set median label) for the node on which the window was mapped. d) Task loss versus the topological loss for SOM-VAE (with Gaussian neighbourhood), DESOM and SOM-CPC (both with and without $sg[\cdot]$). The different curves display various values of α , for which the DESOM model seems most sensitive. The SOM-CPC models follow a smooth optimization curve, minimizing both the task and topological loss, while these losses seem to be more conflicting in SOM-VAE and DESOM training.




Figure 3: Progression of SOM training (nodes in black) in the SOM-CPC model, using either a Gaussian (top) or plus neighbourhood (bottom) kernel, with the PCA projection of the test set latent space plotted behind the nodes. The Gaussian kernel enforces a more strict organization of SOM nodes that quantize the latent space by placing more nodes at higher density areas.

on top of a PCA projection of the continuous test set latents), during training for SOM-CPC with the two type of kernels. It can indeed be seen that the Gaussian kernel enforces a more strict topological organization. Interestingly, the kernel does not only influence the codebook vectors, but also seems to influence the organization of the latent space, seen from the differently-shaped PCA projections in the background.

4.2. Physiological Data

Data: We analyse SOM-CPC on physiological data. To this end, we selected whole-night polysomnography recordings which comprise a broad range of widely-used physiological sensors: electroencephalography (EEG), electromyography (EMG), and electrooculography (EOG). We used subset 3 of the Montreal Archive of Sleep Studies (MASS) database (O’Reilly et al., 2014)¹, consisting of polysomnography recordings for which every 30-second window is labelled

¹<http://ceams-carsm.ca/mass/>

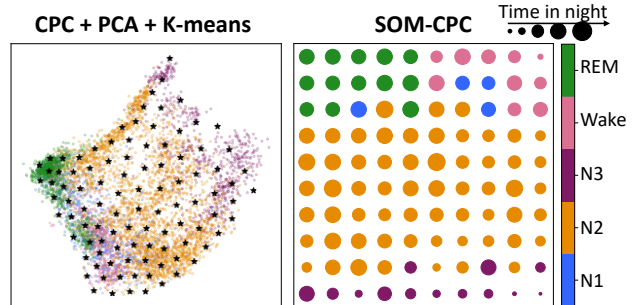


Figure 4: Deep sleep N3 is isolated from light sleep N1, Wake and REM sleep with a cluster of medium-deep sleep N2.

with a sleep stage label from $\{N1, N2, N3, \text{REM}, \text{Wake}\}$. The 62 recordings (from 62 unique subjects) were randomly split into a training ($n = 48$), validation ($n = 8$) and hold-out test set ($n = 7$). Details on the data preprocessing can be found in Appendix A.4.1.

Comparison to deep-SOM baselines: Table 5 in Appendix A.4.3 shows that SOM-CPC again clearly outperformed SOM-VAE and DESOM on all metrics. Using SimCLR (Chen et al., 2020) - a popular contrastive learning framework - as task objective instead of CPC’s InfoNCE loss also deteriorated performance (see Appendix A.4.3).

Ablations: Whether or not the gradients of the SOM loss were stopped towards the encoder did not greatly influence SOM-CPC performance. Topological ordering, measured by TE, and temporal smoothness ($\ell_{2, \text{smooth}}$) deteriorated when changing the Gaussian kernel to a plus kernel, or training CPC and SOM disjointly (Huijben et al., 2022).

Comparison to other baselines: SOM-CPC’s classification performance was higher than that of CPC with $F = 2$, and

CPC followed by PCA. Figure 4 shows the test set PCA projection of the latent space of CPC ($F = 128$), with the K-means nodes as black stars (left), and the SOM (right) trained by the SOM-CPC model (nodes are colored with the most-occurring label in the test set). Both visualizations show similar clustering patterns: deep sleep N3 is isolated from lighter forms of sleep (i.e. N1, Wake and REM sleep) by a thick cluster of medium-deep sleep N2. However, the higher performance of SOM-CPC (see Table 5) indicates that more information is preserved in the 2D space resulting from the SOM-CPC model. The size of the nodes in the SOM of SOM-CPC indicates the average time in the night of windows on that node. A difference is visible in node sizes within the Wake, N2 and N3 clusters, suggesting a possible existence of different sub-categories of sleep within the pre-defined sleep stages.

4.3. Audio

Data: We use a subset of the publicly available LibriSpeech dataset (Panayotov et al., 2015)², which contains multiple minute-long English voice recordings of 251 different speakers, sampled at 16 KHz. We leveraged the publicly available train-test split, as provided by Oord et al. (2019)³, and created an additional validation set by randomly selecting 25% of the training set. Recordings of the ten speakers with the longest recording time were selected to alleviate computational burden. This resulted in a total of 150.9, 54.6, and 46.5 minutes in the training, validation, respectively test set. Model and training details can be found in Appendix A.5.1.

Comparison to deep-SOM baselines: Table 7 in Appendix A.5.2 shows the results of SOM-CPC (which includes a GRU for this dataset), compared to different variants of the DESOM model. SOM-CPC outperforms all DESOM variants by a wide margin and for all choices of the α parameter. The difference in performance between DESOM and SOM-CPC is also visible in Figure 5. SOM-CPC has clustered the SOM nodes belonging to the same speaker, and seems to group male and female speakers (denoted with the node’s shape), while these effects are not present in the SOM of the GRU-DESOM model.

Comparison other baselines: Minimizing the InfoNCE training objective of CPC with $F = 2$ was found challenging for this dataset, which resulted in non-competitive performance of the linear classifier and K-means clustering trained on the 2D latent space. Using PCA for dimensionality reduction of the high-dimensional CPC latent space (with $F = 512$) performed better, but still inferior to SOM-CPC.

²<https://www.openslr.org/12>

³<https://drive.google.com/drive/folders/1BhJ2umKH3whguxMwifaKtSra0TgAbtfb>




Figure 5: SOM-CPC is able to better cluster different speakers than GRU-DESOM. Stars denote women and dots are men.

The SOM of the SOM-CPC model (Figure 5-right) reveals two separate clusters both for speaker 2 (green) and 3 (red). The LibriSpeech corpus contains multiple recordings from each speaker, grouped by (book) chapters from which the speaker was reading. Interestingly, additional analyses revealed that the red and green sub-clusters represented recordings belonging to different chapters: 99.99% of the test-set windows mapped to the upper red sub-cluster belong to the same chapter, while 99.97% of the windows in the lower red sub-cluster belong to another chapter read by this speaker. Similarly, 100.0% of the test set windows in the right green sub-cluster belong to two chapters read by speaker 2, while 98.91% of the windows in the left green sub-cluster belong to another chapter. An auditory inspection revealed that the room acoustics of the recordings belonging to the chapters in different clusters were different, causing changes in the signals which the SOM-CPC model has picked upon. This division between recordings of the same speaker is not visible in the 2D PCA projection of CPC (with $F = 128$), see Figure 8 in Appendix A.5.2.

5. Discussion

We formulated a framework that generalizes deep-SOM algorithms, and proposed a new member of this family, SOM-CPC, for interpretable 2D representation learning of high-rate data streams.

Earlier proposed deep-SOM models used reconstruction objectives. In contrast SOM-CPC leverages contrastive learning and outperformed these models with a wide gap on a variety of metrics. Moreover, it implicitly enforces temporal smoothness, while autoencoder-based models require additional losses and hyperparameter tuning to achieve this. SOM-CPC’s task loss (InfoNCE from Oord et al. (2019)) was found to align better with the topological SOM objective than a reconstruction loss, as already hypothesized by Mrabah et al. (2020). While for some applications CPC could successfully be trained with a 2D latent space directly, optimization was found to be hampered in case of more intricate data spaces. Compared

to vanilla CPC, SOM-CPC enables pattern recognition and knowledge discovery. The SOM objective did not hamper CPC optimization. In the synthetic setup it even had a regularizing effect, resulting in lower regression error than vanilla CPC. The use of a Gaussian neighbourhood kernel, as opposed to a plus kernel, was found to improve the topological ordering in the SOM. No decisive conclusions could be made regarding gradient blocking from the SOM loss towards the encoder parameters. Allowing these gradients to flow did not hurt performance, so for coding simplicity, we would advise to not detach the SOM loss.

The model contains several hyperparameters; the latent space dimension F , the number of past windows L , the number of positive (P) and negative (N) samples, the number of SOM nodes k , the standard deviation of the Gaussian kernel at the end of training, and the loss trade-off parameter α (see Algorithm 1 in Appendix A.1). Effects of setting different values for F and α were already discussed.

The choice for L and P depends on the expected local stationarity of the signals. CPC assumes that the signal behaves stationary, or at least changes less or in a more predictable manner, within the $L + 1 + P$ windows around the anchor/current window, compared to the negative windows that are drawn further away in time. The optimal choice for P and L is thus use-case dependent. In contrast, independent of use case, the higher N , the more tight the bound on mutual information between the context vector and the future latent space, which is maximized by minimizing the InfoNCE loss (Oord et al., 2019).

Choosing a large number of SOM nodes k provides more flexibility for clustering purposes, compared to choosing a low number of nodes. The model can learn to make nearby nodes very similar in case the additional flexibility is not needed, however optimization performance might be hampered when choosing an excessive number of nodes. As such we set k to a sensible, but large-enough, value in each experiment. Lastly, setting $\sigma^{(n_{\max})}$ can be more intricate. We observed instable behaviour in the synthetic experiments when the Gaussian kernel became too narrow towards the end of training. Also, its exact value may affect the trained SOM; a low value provides good flexibility in learning specialized codebook vectors at the risk of instability, while a high value induces a higher correlation between these vectors, at the cost of flexibility.

Setting an appropriate stopping criterion for self-supervised models is debatable. In literature, models with the best test set performance are sometimes reported (He et al., 2020; Fortuin et al., 2019). This is, however, questionable as it may artificially boost reported performance. As such, we created a validation set for model selection in all experiments. Another challenge arises when dealing with aggregated loss functions, since not all losses may smoothly decay and

the weighted summation of losses may result in a different optimal epoch than the sub-losses separately. Besides, classification performance does not necessarily perfectly align with SOM performance or information preservation (see Figure 7 in Appendix A.4.3).

Except from determining a suitable stopping criterion, metric computation in unsupervised learning can also be discussed. Section 3.3 already shortly elucidated upon the way in which we computed classification/regression metrics. Each SOM node was labeled/colored with information from the training set. It can be questioned whether this node labelling should be done on a training or validation set, or directly on the test set for which performance is reported. In earlier times when more conventional (unsupervised) clustering approaches (e.g. K-means) were used, a train/validation/test split was typically not made. As a result, clustering was directly performed on the one and only (test) data set, that was also used to label the clusters/nodes. Moving towards deep learning based approaches where more hyperparameters need to be set and overfitting can become a larger problem, we found it necessary to report on a test set that was not used for labelling the nodes and/or setting hyperparameters. It should thus be taken into account, that direct comparison to results in other deep-clustering works may need a critical eye to see whether similar procedures were used or not.

We believe that SOM-CPC will facilitate knowledge discovery in real-life time series and opens up new research directions for representation learning of time series. Directions include investigation to whether additions like the soft-cluster assignment, cluster hardening loss or a Gaussian latent prior - which have shown to improve the SOM-VAE model (Manduchi et al., 2021) - improve SOM-CPC performance as well. Moreover, the CPC objective assumes slowly-changing data characteristics within the time frame in which positive samples are drawn. A multi-modal variational future prediction could possibly improve performance for data that do not meet this assumption.

Acknowledgements

This work was supported by Onera Health and the project ‘OP-SLEEP’. The project ‘OP-SLEEP’ is made possible by the European Regional Development Fund, in the context of OPZuid.

References

- Berry, R. B., Brooks, R., Gamaldo, C. E., Harding, S. M., Lloyd, R. M., Marcus, C. L., and Vaughn, B. V. The AASM manual for the scoring of sleep and associated events. *Rules, Terminology and Technical Specifications*, Darien, Illinois, American Academy of Sleep Medicine, 176:2012, 2012.
- Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML*, pp. 1597–1607, 2 2020. URL <http://arxiv.org/abs/2002.05709>.
- Cohen, J. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46, 1960.
- Ferles, C., Papanikolaou, Y., and Naidoo, K. J. Denoising Autoencoder Self-Organizing Map (DASOM). *Neural Networks*, 105:112–131, 9 2018. ISSN 18792782. doi: 10.1016/j.neunet.2018.04.016.
- Forest, F., Lebbah, M., Azzag, H., and Lacaille, J. Deep Architectures for Joint Clustering and Visualization with Self-organizing Maps. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, volume 11607 LNAI, pp. 105–116. Springer Verlag, 2019. ISBN 9783030261412. doi: 10.1007/978-3-030-26142-9_10.
- Forest, F., Lebbah, M., Azzag, H., and Lacaille, J. A survey and implementation of performance metrics for self-organized maps. *arXiv preprint arXiv:2011.05847*, 2020.
- Forest, F., Lebbah, M., Azzag, H., and Lacaille, J. Deep embedded self-organizing maps for joint representation learning and topology-preserving clustering. *Neural Computing and Applications*, 33(24):17439–17469, 12 2021. ISSN 14333058. doi: 10.1007/s00521-021-06331-w.
- Fortuin, V., Hüser, M., Locatello, F., Strathmann, H., and Rätsch, G. Som-Vae: Interpretable discrete representation learning on time series. *7th International Conference on Learning Representations, ICLR 2019*, pp. 1–18, 2019.
- 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.
- Hinton, G. E. and Roweis, S. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15, 2002.
- Huijben, I., Nijdam, A., Hermans, L., Overeem, S., van Gilst, M., and van Sloun, R. Self-organizing maps for contrastive embeddings of sleep recordings. In *Proceedings of the 44th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2022.
- Hyvärinen, A. and Pajunen, P. Nonlinear independent component analysis: Existence and uniqueness results. *Neural networks*, 12(3):429–439, 1999.
- Hyvärinen, A., Sasaki, H., and Turner, R. E. Nonlinear ICA Using Auxiliary Variables and Generalized Contrastive Learning. In *The 22nd International Conference on Artificial Intelligence and Statistics*, volume 89, pp. 859–868, 2018. URL <http://arxiv.org/abs/1805.08651>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kiviluoto, K. Topology preservation in self-organizing maps. In *Proceedings of International Conference on Neural Networks (ICNN’96)*, volume 1, pp. 294–299. IEEE, 1996.
- Kohonen, T. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990.
- Le-Khac, P. H., Healy, G., and Smeaton, A. F. Contrastive representation learning: A framework and review. *IEEE Access*, 8:193907–193934, 2020.
- Lee, C. and Schaar, M. V. Temporal phenotyping using deep predictive clustering of disease progression. *37th International Conference on Machine Learning, ICML 2020*, pp. 5723–5733, 2020.
- Lipton, Z. C. The mythos of model interpretability. *Queue*, 16(3):31–57, 2018.
- Madiraju, N. S. Deep Temporal Clustering: Fully Unsupervised Learning of Time-Domain Features. Technical report, Arizona State University, Arizona, 2018.
- Manduchi, L., Hüser, M., Faltys, M., Vogt, J., Rätsch, G., and Fortuin, V. T-DPSOM: An interpretable clustering method for unsupervised learning of patient health states. *Proceedings of the Conference on Health, Inference, and Learning*, 1(1):236–245, 2021. doi: 10.1145/3450439.3451872.

- Mrabah, N., Bouguessa, M., and Ksantini, R. Adversarial deep embedded clustering: on a better trade-off between feature randomness and feature drift. *IEEE Transactions on Knowledge and Data Engineering*, 34(4), 2020.
- Oord, A. v. d., Li, Y., and Oriol, V. Representation Learning with Contrastive Predictive Coding. *arXiv preprint arXiv:1807.03748*, 2019. URL <http://arxiv.org/abs/1807.03748>.
- O'Reilly, C., Gosselin, N., Carrier, J., and Nielsen, T. Montreal Archive of Sleep Studies: an open-access resource for instrument benchmarking and exploratory research. *Journal of Sleep Research*, 23(6):628–635, 12 2014. ISSN 09621105. doi: 10.1111/jsr.12169. URL <http://doi.wiley.com/10.1111/jsr.12169>.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Pesteie, M., Abolmaesumi, P., and Rohling, R. Deep neural maps. In *International Conference on Learning Representations - Workshop track*, 2018.
- Um, T. T., Pfister, F. M., Pichler, D., Endo, S., Lang, M., Hirche, S., Fietzek, U., and Kulić, D. Data augmentation of wearable sensor data for parkinson's disease monitoring using convolutional neural networks. In *Proceedings of the 19th ACM international conference on multimodal interaction*, pp. 216–220, 2017.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural Discrete Representation Learning. In *Conference on Neural Information Processing System*, 2017.
- Woo, G., Liu, C., Sahoo, D., Kumar, A., and Hoi, S. CoST: Contrastive Learning of Disentangled Seasonal-Trend Representations for Time Series Forecasting. In *International Conference on Learning Representations*, 2 2022. URL <http://arxiv.org/abs/2202.01575>.
- Xie, J., Girshick, R., and Farhadi, A. Unsupervised deep embedding for clustering analysis. In *International Conference on Machine Learning*, pp. 478–487. PMLR, 2016.
- Yang, B., Fu, X., Sidiropoulos, N. D., and Hong, M. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 3861–3870. PMLR, 06–11 Aug 2017.
- Yin, H. The self-organizing maps: background, theories, extensions and applications. *Computational intelligence: A compendium*, pp. 715–762, 2008.
- Zimmermann, R. S., Sharma, Y., Schneider, S., Bethge, M., and Brendel, W. Contrastive Learning Inverts the Data Generating Process. In *International Conference on Machine Learning*, pp. 12979–12990, 2021. URL <http://arxiv.org/abs/2102.08850>.

A. Experimental details

This appendix contains all information for full reproducibility of the experiments. Domain-independent details of SOM-CPC and its evaluation are provided in Appendix A.1, while benchmark implementations are discussed in Appendix A.2. Domain-specific settings for the experiments on synthetic data, physiological data, and audio experiments are discussed in sections A.3, A.4, and A.5, respectively.

A.1. General details

Algorithm 1 provides pseudocode of SOM-CPC, when considering the presence of an AR module. For reproducibility, the full code base can be found at <https://github.com/IamHuijben/SOM-CPC.git>.

Algorithm 1 SOM-CPC

Input: Dataset \mathcal{X} , model comprising $f_\theta, g_\psi, \{\mathbf{W}_p\}_{p=1}^P, q_\Phi$, latent space dimension F , # past windows L , # positive samples P , # negative samples N , # SOM nodes k , Gaussian neighbourhood function \mathcal{S} with $\sigma^{(n_{\max})}, n_{\max}$, loss trade-off parameter α .

Output: A trained SOM: topologically ordered codebook Φ that represents data \mathcal{X} in 2D.

```

- Randomly initialize  $\Phi \sim \text{U}[-\sqrt{1/F}, \sqrt{1/F}]$ 
- Initialize the Gaussian kernel according to Equation (2) with:  $\sigma^{(0)} = \frac{1}{2}\sqrt{k}$  and  $\lambda = -n_{\max}/\log(\sigma^{(n_{\max})}/\sigma^{(0)})$ 
for  $n$  in  $n_{\max}$  do
  for batch in # batches do
    - Sample a sequence of datapoints:  $[\mathbf{x}(t-L), \dots, \mathbf{x}(t)] \subset \mathcal{X}$ 
    - Define  $P$  positive samples:  $\{\mathbf{x}(t+p)\}_{p=1}^P$ 
    - Sample  $P \times N$  negative samples  $\mathcal{X}' \subset \mathcal{X}$ , with  $|\mathcal{X}'| = P \times N$ 
    - Encode:
      - data sequence:  $\mathbf{c}(t) = g_\psi(f_\theta([\mathbf{x}(t-L), \dots, \mathbf{x}(t)]))$ 
      - positive samples:  $\{\mathbf{z}(t+p)\}_{p=1}^P = f_\theta(\{\mathbf{x}(t+p)\}_{p=1}^P)$ 
      - negative samples:  $\mathcal{Z}' = f_\theta(\mathcal{X}')$ 
      - Compute  $\mathcal{L}_{\text{topo}}(\mathbf{c}(t))$  and  $\mathcal{L}_{\text{infoNCE}}$  according to Equation (6) and Equation (9)
      - Update trainable parameters  $\propto \alpha \mathcal{L}_{\text{topo}}$  and  $\mathcal{L}_{\text{infoNCE}}$ 
    end for
  - Update  $\sigma^{(n)}$  according to Equation (4)
end for

```

Several metrics were used to quantify SOM performance, as explained in Section 3.3. The purity implementation was taken from the Github implementation of Fortuin et al. (2019), while NMI was computed using the sklearn library (Pedregosa et al., 2011). The topographic error implementation comes from the *SOMperf* python library (Forest et al., 2020). Finally, the $\ell_{2,\text{smooth}}$ distance is computed using the *norm* function in the *linalg* library of Numpy.

We here provide the formal definitions of all reported metrics. After SOM training, its nodes are ‘colored’ using a (sub)set of data, which assigns labels ν_i to each node ϕ_i (with $i \in \{1, \dots, k\}$). When running inference on a data stream \mathbf{X} and corresponding labels \mathbf{Y} , each window \mathbf{x} in \mathbf{X} is mapped to a node/codebook vector. If \mathbf{x} is mapped to node j , then the predicted label for \mathbf{x} equals $\hat{y} = \nu_j$. $\hat{\mathbf{Y}}$ contains the predicted labels for all data points in \mathbf{X} . The definition of all metrics, evaluated over data (\mathbf{X}, \mathbf{Y}) , are provided below and can straightforwardly be applied for the full dataset $(\mathcal{X}, \mathcal{Y})$ by averaging the metrics for each (\mathbf{X}, \mathbf{Y}) . Purity is defined as:

$$\text{purity} = \frac{1}{|\mathbf{X}|} \sum_{i=1}^k \sum_{\mathbf{x}, y \in \mathbf{X}, \mathbf{Y}} \mathbb{1}(q_\Phi(\mathbf{c}) = \phi_i, y = \nu_i),$$

with $\mathbb{1}$ an indicator function that returns one if the condition is met, and zero otherwise. Normalized Mutual Information is defined as:

$$\text{NMI} = \frac{\text{MI}(\mathbf{Y}; \hat{\mathbf{Y}})}{\sqrt{\text{H}(\mathbf{Y}) \times \text{H}(\hat{\mathbf{Y}})}},$$

with MI being the Mutual Information and H the Shannon entropy. The regression and classification metrics are:

$$\text{SE}_{\text{target}} = \frac{1}{|\mathbf{X}|} \sum_{x,y \in \mathbf{X}, \mathbf{Y}} (\hat{y} - y)^2,$$

$$\text{Cohen's Kappa} = \frac{p_a(\mathbf{Y}, \hat{\mathbf{Y}}) - p_c(\mathbf{Y}, \hat{\mathbf{Y}})}{1 - p_c(\mathbf{Y}, \hat{\mathbf{Y}})},$$

where p_a and p_c denote the probability of agreement, and the probability of agreement by chance, respectively. Lastly, metric definitions concerning topological organization are:

$$\ell_{2,\text{smooth}} = \frac{1}{|\mathbf{X}|} \sum_{x \in \mathbf{X}} \left\| \mathcal{P}[q_\Phi(\mathbf{c}(t))], \mathcal{P}[q_\Phi(\mathbf{c}(t+1))] \right\|_2,$$

$$\text{TE} = \frac{1}{|\mathbf{X}|} \sum_{x \in \mathbf{X}} \mathbb{1} \left(\left\| \mathcal{P}[q_\Phi(\mathbf{c})], \mathcal{P}[\tilde{q}_\Phi(\mathbf{c})] \right\|_2 > 1 \right),$$

where \mathcal{P} projects a codebook vector to its corresponding coordinate on the 2D grid, and $q_\Phi(\mathbf{c})$ and $\tilde{q}_\Phi(\mathbf{c})$ denote the winning, respectively, second-best codebook vector for embedding \mathbf{c} . Note that in case the AR module is not used in SOM-CPC all metrics are computed on \mathbf{z} rather than \mathbf{c} . To enhance readability, time indices (t) were omitted for metrics that do not depend on windows from varying times.

A.2. Benchmarks and ablations

To benchmark our implementation of the SOM-VAE model (and the very similar DESOM model that used the same backbone implementation in our code base), we replicated the results on MNIST. MNIST was not used for further experimentation with SOM-CPC in the main body of this paper since this work focuses on high-rate time series. Using $k = 16$ SOM nodes, Fortuin et al. (2019) report purity = 0.731 ± 0.004 and NMI = 0.594 ± 0.004 in table 1, and purity = 0.721 ± 0.006 and NMI = 0.587 ± 0.003 in table S1. All numbers are averages and standard errors over 10 runs.

Settings that were not provided in the paper, were taken from the hard-coded settings that we found in the provided code base. Instead of splitting the standard MNIST training set in a training and test split, as done by Fortuin et al. (2019), we used the available train/test split that comes with the standard MNIST dataloader from Pytorch. From the first ten runs, one run fully collapsed and resulted in extremely poor performance. Considering this run as an outlier, we run an 11th run and here report the average and standard errors of the 10 non-collapsed runs: purity = 0.705 ± 0.002 and NMI = 0.584 ± 0.001 . Our performance does reasonably well match the reported performance by Fortuin et al. (2019).

To restrict the search space of hyperparameters in the SOM-VAE(-prob) model, which has multiple loss multipliers, we fixed some of these multipliers for further experiments in this paper. Multiplier β scales the SOM loss that sums the quantization error of the four neighbour nodes in the plus kernel. It is, therefore, expected to be at least 4 times larger than the commitment loss, which only reflects the quantization error of the winning node. Choosing $\beta = \alpha/4$ would imply that the weighing of the summed quantization error of the four neighbours is equal to the weighing of this error for the winning node. To give the winning node a slightly higher importance, we set $\beta = \alpha/5 = 0.2\alpha$. The authors of SOM-VAE (Fortuin et al., 2019) used, instead, a search strategy to find the optimal setting. Their code base⁴ shows that the SOM loss was multiplied with 0.9, while $\alpha = 1$. Taking into account that their implementation of the commitment loss averaged the quantization error of the four neighbour nodes, while we summed the contribution, their effective setting was thus set to $\beta = \frac{0.9}{4}\alpha = 0.225\alpha$, which is close to what we used in our experiments.

⁴https://github.com/ratschlab/SOM-VAE/blob/master/som_vae/somvae_train.py, line 79.

We compared SOM-CPC also against vanilla CPC training followed by a linear classifier or K-means clustering, while freezing the encoder parameters. The supervised linear classifier took in all experiments the form of one fully-connected layer, including biases, that was followed by a log-softmax activation for the physiological and audio cases. It was trained using the mean squared error for the synthetic data set, and cross-entropy loss for physiological and audio experiments. K-means clustering was run from the sklearn library, with the default settings. The number of clusters was chosen to be equal to the number of nodes in the SOM-CPC models against which the performance was compared. Also the disjointly-trained SOM had exactly the same settings as the SOM in the SOM-CPC models with which it was compared. Disjoint training was implemented as follows. First the vanilla CPC model (with $F = 128$) was trained, after which the training data was projected into the latent space using the trained encoder. Then the SOM was trained using the loss function in Equation (6), while using the fixed latent representations as input.

Several ablations were performed on the SOM-CPC model. We tested the effect of propagating gradients of \mathcal{L}_{SOM} to the encoder parameters, the difference between using a Gaussian neighbourhood kernel versus a plus kernel, and the effect of jointly training CPC and the SOM. Moreover, the effect of certain settings that are typically used in the CPC objective are investigated in this appendix. CPC’s InfoNCE objective for one window p , given in Equation (10), can as follows be generalized to a more general contrastive learning objective:

$$\mathcal{L}_p = -\mathbb{E}_{\mathcal{X}} \left[\log \frac{\exp(\text{sim}(z_a, z_p)/\tau)}{\sum_{z' \in \mathcal{Z}'_p \cup \{z_p\}} \exp(\text{sim}(z_a, z')/\tau)} \right], \tag{11}$$

where z_a is the latent space of the current (or *anchor*) window, $\text{sim}(\cdot)$ a similarity metric, τ a temperature parameter and the other symbols are equivalent to Equation (10). CPC typically uses $\tau = 1$, and the dot product as the similarity metric. However, other related contrastive learning objectives, e.g. in SimCLR (Chen et al., 2020), use a temperature value that is often set to 0.07 (Chen et al., 2020; Woo et al., 2022), and a cosine similarity instead of the (unnormalized) dot product. As such, in all experiments we added ablations where we set τ at 1 or 0.07, and use either the dot-product or the cosine similarity as the similarity metric. Results are discussed in Appendices A.3.2, A.4.3 and A.5.2.

A.3. Synthetic experiments

A.3.1. TRAINING DETAILS

Table 2 summarizes the encoder and decoder architectures used in this experiment. The *output size* column in the table uses channels-first notation. The SOM-VAE and DESOM models were found to benefit from a convolutional part of the encoder that did not fully reduce the temporal dimension to size 1. As such, the last convolutional layer of the SOM-CPC encoder was changed for a fully connected layer preceded by a flattening operation for the autoencoder-based models. The SOM-CPC and CPC model are run without an AR module, to make the fairest comparison to the SOM-VAE(prob) and DESOM models, which also do not incorporate such a component.

For SOM-CPC, $P = 3$ future predictions (i.e. positive samples) were used, and $N = 3$ negative samples were drawn for each positive sample. The latter were drawn randomly from the entire training set. The standard deviation of the Gaussian neighbourhood kernel was exponentially decayed until $\sigma^{(r_{\max})} = 2$. Choosing a lower value at the end of training induced instable optimization behaviour.

All models (including the benchmarks) were trained using the Adam optimizer (Kingma & Ba, 2015), with a learning rate of 0.001 and a batch size of 128. Each model was trained for maximally 1000 epochs. The best model was selected based on the lowest task loss on the validation set, being $\mathcal{L}_{\text{recon}}$ for SOM-VAE and DESOM and $\mathcal{L}_{\text{InfoNCE}}$ for SOM-CPC and CPC. We did not use the full training objective $\mathcal{L}_{\text{deep-SOM}}$ as model selection criterion, as both the commitment and SOM loss showed to be low initially (possibly due to low values of the random initialization of the model), while both increased and reached a steady-state later in training. The linear classifier and the disjointly-trained SOM on the CPC embeddings were trained until convergence, for maximally 1000 epochs.

A.3.2. EXTENDED RESULTS

Table 3 extends Table 1 with additional sweeps of hyperparameters α , γ , and ζ , and ablations with different settings of the temperature value τ and the used similarity metric. Compared to SOM-VAE, SOM-VAE-prob is expected to show more smooth trajectories over time, captured in the $\ell_{2,\text{smooth}}$ distance metric, thanks to the additional transition and smoothness

Table 2: Model details for the synthetic data experiments in Section 4.1.

Layer type	Output size	Channels	Activation	Kernel size	Strides	Dilation	Padding
Encoder for SOM-CPC and CPC							
Conv1D	bs \times 16 \times 128	16	Leaky ReLU (0.01)	9	1	1	same
MaxPool1D	bs \times 16 \times 32	-	-	4	4	-	-
Dropout (0.1)	bs \times 16 \times 32	-	-	-	-	-	-
Conv1D	bs \times 32 \times 32	32	Leaky ReLU (0.01)	7	1	1	same
MaxPool1D	bs \times 32 \times 8	-	-	4	4	-	-
Dropout (0.1)	bs \times 32 \times 8	-	-	-	-	-	-
Conv1D	bs \times 64 \times 8	64	Leaky ReLU (0.01)	3	1	1	same
MaxPool1D	bs \times 64 \times 2	-	-	4	4	-	-
Dropout (0.1)	bs \times 64 \times 2	-	-	-	-	-	-
Conv1D	bs \times 128 \times 2	128	Leaky ReLU (0.01)	3	1	1	same
MaxPool1D	bs \times 128 \times 1	-	-	2	2	-	-
Encoder for SOM-VAE(-prob) and DESOM							
Conv1D	bs \times 16 \times 128	16	Leaky ReLU (0.01)	9	1	1	same
MaxPool1D	bs \times 16 \times 32	-	-	4	4	-	-
Dropout (0.1)	bs \times 16 \times 32	-	-	-	-	-	-
Conv1D	bs \times 32 \times 32	32	Leaky ReLU (0.01)	7	1	1	same
MaxPool1D	bs \times 32 \times 8	-	-	4	4	-	-
Dropout (0.1)	bs \times 32 \times 8	-	-	-	-	-	-
Conv1D	bs \times 64 \times 8	64	Leaky ReLU (0.01)	3	1	1	same
Flatten	bs \times 512	-	-	-	-	-	-
Fully Connected	bs \times 128	128	Leaky ReLU (0.01)	-	-	-	-
Decoder for SOM-VAE(-prob) and DESOM							
Fully Connected	bs \times 512	512	Leaky ReLU (0.01)	-	-	-	-
Unflatten	bs \times 64 \times 8	-	-	-	-	-	-
Conv1D	bs \times 32 \times 8	32	Leaky ReLU (0.01)	3	1	1	same
ConvTranspose1D	bs \times 32 \times 32	32	None	4	4	1	0
Conv1D	bs \times 16 \times 32	16	Leaky ReLU (0.01)	7	1	1	same
ConvTranspose1D	bs \times 16 \times 128	16	None	4	4	1	0
Conv1D	bs \times 1 \times 128	1	Tanh	9	1	1	same

loss (multiplied by γ and ζ , respectively). It can be seen that tuning these hyperparameters is a complex process, and a sweep did not result in one SOM-VAE-prob run that performed better than the best SOM-VAE model. In contrary, the addition of the extra losses possibly interfered with the optimization process, and only very delicate settings of γ and ζ might improve model performance eventually. A sweep over the topological loss multiplier α , revealed a low sensitivity of SOM-CPC to this value. It can be seen that changing the temperature value and/or the similarity metric did not significantly alter the performance consistently on all reported metrics.

Figure 6 shows PCA projections of the latent space of the SOM-VAE, DESOM, and SOM-CPC models that are indicated with a * in Tables 1 and 3, and for which the SOMs were visualized in Figure 2. Organization of the signal frequencies is much better for the SOM-CPC model, providing an explanation for the better (i.e. lower) SE_{target} of this model, as compared to SOM-VAE and DESOM.

Table 3: This is the extended version of Table 1 on the synthetic data results, including a sweep of hyperparameters α , γ and ζ . Bold values indicate the best performance per column (excluding the upper bound of the vanilla CPC model, which does not result in a 2D representation). The models indicated with a * were used to depict trained SOMs in Figure 2 and PCA projections in Figure 6.

Model	α	S	$\mathcal{L}_{\text{SOM}}^{\text{sg}}[-]$	$\text{SE}_{\text{target}}$	$\ell_{2,\text{smooth}}$	TE	
CPC + linear classifier	-	-	-	2.62± 2.37	-	-	
CPC + K-means	-	-	-	1.09± .62	-	-	
CPC ($F = 2$) + linear classifier	-	-	-	25.01±42.94	-	-	
CPC ($F = 2$) + K-means	-	-	-	.76± 1.31	-	-	
CPC + PCA + linear classifier	-	-	-	42.81±58.12	-	-	
CPC + PCA + K-means	-	-	-	4.42± 9.01	-	-	
SOM-VAE	1e-3	Plus	✓	11.59±13.69	2.60±.46	.38±.04	
	1e-2	Plus	✓	14.57±44.10	2.98±.84	.38±.05	
*	.1	Plus	✓	8.02± 4.58	2.41±.68	.28±.07	
	1	Plus	✓	13.43± 3.66	2.75±.45	.33±.05	
SOM-VAE	1e-5	Gaussian	✓	11.12±17.05	1.93±.29	.07±.03	
	1e-4	Gaussian	✓	1.52±26.61	1.95±.36	.08±.03	
	1e-3	Gaussian	✓	11.60±25.43	1.92±.34	.06±.02	
	1e-2	Gaussian	✓	18.13±48.66	2.03±.42	.09±.03	
SOM-VAE-prob	($\gamma = 5e-5, \zeta = 1e-3$)	1e-3	Plus	✓	21.26±55.00	3.78±.52	.93±.04
	($\gamma = 4e-5, \zeta = 1e-3$)	1e-3	Plus	✓	21.30±37.37	4.23±.57	.88±.05
	($\gamma = 3.3e-5, \zeta = 1e-3$)	1e-3	Plus	✓	22.52±48.97	3.81±.51	.98±.02
	($\gamma = 5e-4, \zeta = 1e-2$)	1e-2	Plus	✓	14.65±19.58	3.70±.51	.86±.08
	($\gamma = 4e-4, \zeta = 1e-2$)	1e-2	Plus	✓	27.25±72.82	3.54±.67	.94±.02
	($\gamma = 3.3e-4, \zeta = 1e-2$)	1e-2	Plus	✓	21.62±38.78	3.71±.76	.97±.01
	($\gamma = 5e-5, \zeta = 1e-2$)	.1	Plus	✓	26.82±68.84	3.23±.80	.68±.07
	($\gamma = 4e-5, \zeta = 1e-2$)	.1	Plus	✓	22.34±64.92	3.11±.73	.74±.07
	($\gamma = 3.3e-5, \zeta = 1e-2$)	.1	Plus	✓	2.10±48.80	3.15±.73	.63±.05
	($\gamma = 5e-4, \zeta = .1$)	.1	Plus	✓	4.85±75.96	3.06±.55	.84±.05
	($\gamma = 4e-4, \zeta = .1$)	.1	Plus	✓	29.96±46.96	2.81±.34	.82±.13
	($\gamma = 3.3e-4, \zeta = .1$)	.1	Plus	✓	3.96±56.94	3.95±.83	.85±.11
DESOM	1e-5	Gaussian	✗	14.00± 3.10	1.99±.36	.13±.07	
	1e-4	Gaussian	✗	19.09±44.04	1.95±.28	.11±.05	
	1e-3	Gaussian	✗	12.58±27.10	1.92±.31	.08±.03	
	1e-2	Gaussian	✗	13.66±44.71	1.89±.33	.07±.03	
*	.1	Gaussian	✗	10.77±10.85	2.20±.44	.06±.02	
	1	Gaussian	✗	22.86±55.71	2.26±.43	.09±.03	
SOM-CPC (ours)	1e-5	Gaussian	✗	.95± 2.40	1.24±.31	.05±.02	
*	1e-4	Gaussian	✗	.72± 1.08	1.37±.37	.02±.01	
	1e-3	Gaussian	✗	.81± .75	1.06±.28	.03±.01	
	1e-2	Gaussian	✗	.62± .71	1.08±.28	.06±.04	
	.1	Gaussian	✗	1.90± 4.07	1.18±.30	.04±.02	
	1e-5	Gaussian	✓	.64± .71	1.04±.26	.05±.02	
	1e-4	Gaussian	✓	.68± .67	1.19±.43	.06±.03	
	1e-3	Gaussian	✓	.75± 1.02	1.12±.32	.06±.03	
	1e-2	Gaussian	✓	.47± .48	.99±.24	.07±.04	
	.1	Gaussian	✓	.89± 1.50	1.16±.32	.07±.03	
	1e-4	Plus	✗	1.71± 1.15	2.46±.51	.30±.06	
	1e-2	Plus	✓	1.16± .61	1.85±.26	.12±.04	
SOM-CPC ($\tau = 0.07, \text{sim} = \text{cosine sim.}$)	1e-4	Gaussian	✗	1.47± 2.60	1.15±.34	.01±.02	
SOM-CPC ($\tau = 1, \text{sim} = \text{cosine sim.}$)	1e-4	Gaussian	✗	2.26± 4.91	.96±.13	.06±.02	
SOM-CPC ($\tau = 0.07, \text{sim} = \text{dot prod.}$)	1e-4	Gaussian	✗	1.22± 4.08	1.15±.37	.07±.05	
CPC + SOM (disjoint)	-	Gaussian	-	.84± 1.14	1.47±.50	.03±.01	

Ablations




Figure 6: PCA projections of the continuous latent spaces of the full test set for the SOM-VAE, DESOM, and SOM-CPC models of which the SOMs were visualized in Figure 2. Organization of the signal frequencies is most structured in the SOM-CPC model.

A.4. Physiological data experiments

A.4.1. DATA PROCESSING

For the experiments on physiological data, we used subset 3 of the publicly available Montreal Archive of Sleep Studies (MASS) database (O’Reilly et al., 2014), consisting of 62 whole-night polysomnography recordings. Each recording contains, among others, electroencephalography (EEG), chin electromyography (EMG), and electrooculography (EOG) data. We refer the reader to (O’Reilly et al., 2014) for more details regarding this dataset. We selected the channels that are typically used in clinical practice, comprising three EEG channels (F4, C4, O2), the two EOG channels, and one chin EMG derivation, and downsampled the data to 128 Hz. Sleep stage labels that follow the guidelines of the American Academy of Sleep Medicine (AASM) (Berry et al., 2012) (being wakefulness (Wake), rapid-eye movement (REM) sleep, or non-REM1 till non-REM3 (N1, N2, N3)) were available for every non-overlapping 30-second window, the common label resolution in clinical practice.

Some processing had already been done by the distributors of the MASS dataset (O’Reilly et al., 2014). The 60 Hz powerline interference was, however, not fully suppressed, and we wanted to down sample each signal to 128 Hz to reduce computational complexity. As such, before downsampling, all derivations were additionally filtered with a zero-phase (i.e. two-directional) 5th order Butterworth band-pass filter (0.3 – 59 Hz), followed by another zero-phase 5th order Butterworth notch filter (59 – 61 Hz). Channels were normalized within-patient and per channel, yielding mean subtraction, followed by normalization such that amplitudes of 95% of the samples were mapped between -1 and +1. The 62 recordings (numbered 1 – 64, with number 43 and 49 missing) were split into a training set including patients 1 – 48 ($n = 47$), a validation set including patients 50 – 57 ($n = 8$), and hold-out test set that included patients 58 – 64 ($n = 7$).

A.4.2. TRAINING DETAILS

Dimensionality reduction of polysomnography data was done by encoding all selected channels in each non-overlapping 30-second window using standard convolutional encoder. Table 4 summarizes the used encoder and decoder (for SOM-VAE and DESOM) architectures. The latent space for decoding in the SOM-VAE and DESOM benchmark models was not fully reduced to a 1D vector to enhance training. Nevertheless, the last adaptive average pooling layer that was used in the encoder of SOM-CPC, was applied in the bottleneck of SOM-VAE and DESOM before SOM quantization took place. As a result the codebook vectors in all models were of size $F = 128$. The decoder architecture (see Table 4) was used both for the continuous and discrete decoding in the SOM-VAE model (without weight tying). No AR-component was used in the SOM-CPC model to make a fair comparison to the SOM-VAE and DESOM model that also did not include such a component.

For the SOM-CPC model, $P = 3$ future predictions (i.e. positive samples) were used, and $N = 3$ negative samples were drawn for each positive sample. The latter were drawn from the same subject as the positive sample. The σ of the Gaussian neighbourhood kernel was exponentially annealed to $\sigma^{(n_{\max})} = 0.5$ during training.

All models were trained with the Adam optimizer (Kingma & Ba, 2015), with a learning rate of 1e-4 and a batch size of 128. Each model was trained for maximally 500 epochs, and the best model was selected based on the lowest $\mathcal{L}_{\text{recon}}$ (for SOM-VAE and DESOM) or $\mathcal{L}_{\text{InfoNCE}}$ (for SOM-CPC and CPC) on the validation set.

A.4.3. EXTENDED RESULTS

Table 5 shows the quantitative results on physiological data, comparing SOM-CPC against deep-SOM models (SOM-VAE and DESOM) and disjoint training of CPC, followed by either a supervised linear classifier, K-means or a SOM. Discussion of the main results in this table can be found in Section 4.2. The ablation experiments in which the value of τ and/or the similarity metric was altered show that classification and clustering performance slightly dropped when using the cosine similarity with a temperature value of 1, while the topographic organization slightly improved (i.e. lower TE). These effects vanished when using a temperature of $\tau = 0.07$. Both runs showed worse temporal smoothness (i.e. higher $\ell_{2,\text{smooth}}$). As expected, only changing the temperature value to 0.07 did almost not affect results, suggesting that the linear projector heads were able to adjust for this scaling factor.

We also tested the performance when using the SimCLR (Chen et al., 2020) objective for the task loss, instead of the CPC objective. SimCLR is also a contrastive learning framework, but instead of drawing positive samples from the future latent space, these samples are created by applying augmentations on the anchor window. Inspired by (Um et al., 2017) we used the following augmentations: independent and identically distributed Gaussian noise $\mathcal{N}(0, 0.05)$ was added (called jitter in

Table 4: Model details for the experiments on physiological data in Section 4.2.

Layer type	Output size	Channels	Activation	Kernel size	Strides	Dilation	Padding
Encoder for SOM-CPC and CPC							
Conv1D	bs \times 16 \times 3826	16	Leaky ReLU (0.01)	15	1	1	0
MaxPool1D	bs \times 16 \times 765	-	-	5	5	-	-
Dropout (0.1)	bs \times 16 \times 765	-	-	-	-	-	-
Conv1D	bs \times 32 \times 757	32	Leaky ReLU (0.01)	9	1	1	0
MaxPool1D	bs \times 32 \times 151	-	-	5	5	-	-
Dropout (0.1)	bs \times 32 \times 151	-	-	-	-	-	-
Conv1D	bs \times 64 \times 147	64	Leaky ReLU (0.01)	5	1	1	0
MaxPool1D	bs \times 64 \times 29	-	-	5	5	-	-
Dropout (0.1)	bs \times 64 \times 29	-	-	-	-	-	-
Conv1D	bs \times 128 \times 27	128	Leaky ReLU (0.01)	3	1	1	0
AdaptiveAvgPool1D	bs \times 128 \times 1	-	-	-	-	-	-
Encoder for SOM-VAE and DESOM							
Conv1D	bs \times 16 \times 3861	16	Leaky ReLU (0.01)	15	1	1	(18, 17)
MaxPool1D	bs \times 16 \times 772	-	-	5	5	-	-
Dropout (0.1)	bs \times 16 \times 772	-	-	-	-	-	-
Conv1D	bs \times 32 \times 764	32	Leaky ReLU (0.01)	9	1	1	0
MaxPool1D	bs \times 32 \times 152	-	-	5	5	-	-
Dropout (0.1)	bs \times 32 \times 152	-	-	-	-	-	-
Conv1D	bs \times 64 \times 148	64	Leaky ReLU (0.01)	5	1	1	0
MaxPool1D	bs \times 64 \times 29	-	-	5	5	-	-
Dropout (0.1)	bs \times 64 \times 29	-	-	-	-	-	-
Conv1D	bs \times 128 \times 27	128	Leaky ReLU (0.01)	3	1	1	0
Decoder for SOM-VAE and DESOM							
Conv1D	bs \times 64 \times 27	64	Leaky ReLU (0.01)	3	1	1	1
ConvTranspose1D	bs \times 64 \times 135	64	None	5	5	1	0
Conv1D	bs \times 32 \times 135	32	Leaky ReLU (0.01)	5	1	1	2
ConvTranspose1D	bs \times 32 \times 675	32	None	5	5	1	0
Conv1D	bs \times 16 \times 675	16	Leaky ReLU (0.01)	9	1	1	4
ConvTranspose1D	bs \times 16 \times 3375	16	None	5	5	1	0
Conv1D	bs \times 6 \times 3375	6	None	15	1	1	7
Crop	bs \times 6 \times 3340	-	-	-	-	-	-

their implementation), each channel was scaled with a value drawn from $\mathcal{N}(0, 0.1)$, windows were split in 4 sub-windows of minimal 2 seconds and randomly permuted, and lastly time series were both time warped and magnitude warped. The latter two augmentations make use of smooth curves that smoothly vary the positions of time stamps or magnitude values, respectively.

Besides the difference on how to create positive samples, the originally proposed SimCLR model has some other slight differences with respect to the CPC model:

- The SimCLR loss uses the cosine similarity, while CPC uses the (unnormalized) dot product as the similarity metric (see Equation (11)).
- SimCLR uses an additional temperature τ in its loss function (see Equation (11)), for which the value is often set to 0.07 (Chen et al., 2020; Woo et al., 2022). CPC does not incorporate such a temperature, which effectively means that it uses a value of 1.
- SimCLR uses a non-linear MLP projection head, while CPC uses linear projection heads.
- SimCLR uses negative samples from within the batch, while this is not specified in the CPC paper. This specified design choice makes SimCLR typically very sensitive to the batch size.
- SimCLR was not proposed to include an auto-regressive component, and can not straightforwardly be extended to do so, while CPC can be implemented with or without such a module.

For the most fair comparison, the procedure for drawing negative samples in SimCLR is done equivalently as for SOM-CPC, i.e. within the recording, instead of within the batch. However, in the SOM-SimCLR model (i.e. the joint training of SOM with SimCLR), each drawn negative sample is added to the set of negative samples both in its raw form, and with a random augmentation, which effectively doubles the number of negative samples. Table 5 reports the performance of SOM-SimCLR (i.e. jointly optimizing SimCLR for feature extraction and a SOM), both for $\tau = 0.07$ and $\tau = 1$, while using the cosine similarity. All settings regarding training procedure and the SOM were set equivalently as in the SOM-CPC training. Table 5 shows that SOM-SimCLR results for $\tau = 0.07$ are better than those with $\tau = 1$, which is in line with findings from (Chen et al., 2020; Woo et al., 2022). However, even with $\tau = 0.07$, performance of SOM-SimCLR is lower on all metrics compared to the SOM-CPC model with the same value for α . The higher $\ell_{2,\text{smooth}}$ metric of SOM-SimCLR

SOM-CPC: Unsupervised Contrastive Learning with Self-Organizing Maps

Table 5: Test set performance of various models trained on physiological recordings. SOMs of models with a * are visualized in Figure 4. Bold values indicate the best performance per column (excluding the upper bound of the vanilla CPC model, which does not result in a 2D representation).

Model	α	S	$\mathcal{L}_{\text{SOM}} \text{sg}[\cdot]$	Purity	NMI	Cohen's kappa	$\ell_{2, \text{smooth}}$	TE
CPC + linear classifier	-	-	-	-	-	.68±.10	-	-
CPC + K-means	-	-	-	.79	.29	.61±.11	-	-
CPC ($F = 2$) + linear classifier	-	-	-	-	-	.52±.10	-	-
CPC ($F = 2$) + K-means	-	-	-	.74	.24	.55±.09	-	-
CPC + PCA + linear classifier	-	-	-	-	-	.54±.09	-	-
* CPC + PCA + K-means	-	-	-	.77	.26	.57±.08	-	-
SOM-VAE	1e-3	Plus	✓	.71	.23	.51±.04	2.36±.26	.24±.03
	1e-2	Plus	✓	.71	.23	.51±.04	2.67±.17	.30±.04
	.1	Plus	✓	.72	.23	.52±.03	2.60±.34	.28±.04
	1	Plus	✓	.71	.23	.53±.03	3.08±.32	.31±.05
DESOM	1e-6	Gaussian	✗	.70	.27	.53±.05	2.14±.32	.10±.02
	1e-5	Gaussian	✗	.70	.23	.50±.04	2.10±.26	.11±.03
	1e-4	Gaussian	✗	.71	.22	.51±.04	2.35±.24	.17±.04
	1e-3	Gaussian	✗	.71	.22	.51±.05	2.40±.16	.22±.01
	1e-2	Gaussian	✗	.71	.22	.50±.04	2.30±.26	.23±.02
SOM-SimCLR ($\tau = 0.07$)	1e-3	Gaussian	✗	.73	.23	.53±.13	2.21±.35	.29±.03
SOM-SimCLR ($\tau = 1$)	1e-3	Gaussian	✗	.70	.20	.48±.16	1.87±.30	.50±.07
SOM-CPC (ours)	1e-5	Gaussian	✗	.78	.27	.59±.11	1.03±.11	.04±.01
	1e-4	Gaussian	✗	.78	.27	.61±.10	1.01±.10	.06±.02
	1e-3	Gaussian	✗	.78	.27	.61±.12	1.02±.09	.03±.01
	1e-2	Gaussian	✗	.79	.28	.60±.11	1.08±.11	.19±.04
	.1	Gaussian	✗	.79	.28	.65±.07	1.09±.09	.19±.04
	1e-3	Gaussian	✓	.78	.27	.62±.10	1.02±.12	.07±.03
	1e-2	Gaussian	✓	.78	.27	.60±.10	1.07±.09	.17±.04
	.1	Gaussian	✓	.78	.27	.62±.10	1.06±.11	.08±.02
	1	Gaussian	✓	.78	.27	.59±.12	1.04±.12	.08±.02
	1e-3	Plus	✗	.79	.28	.61±.11	1.35±.24	.26±.08
	1e-3	Plus	✓	.79	.28	.60±.10	1.38±.28	.27±.08
SOM-CPC ($\tau = 0.07$, sim = cosine sim.)	1e-3	Gaussian	✗	.78	.27	.60±.12	1.36±.13	.07±.02
SOM-CPC ($\tau = 1$, sim = cosine sim.)	1e-3	Gaussian	✗	.73	.27	.58±.11	1.43±.15	.03±.01
SOM-CPC ($\tau = 0.07$, sim = dot prod.)	1e-3	Gaussian	✗	.79	.28	.62±.10	1.06±.11	.06±.02
CPC + SOM (disjoint)	-	Gaussian	-	.79	.28	.62±.11	1.21±.11	.52±.04

Ablations

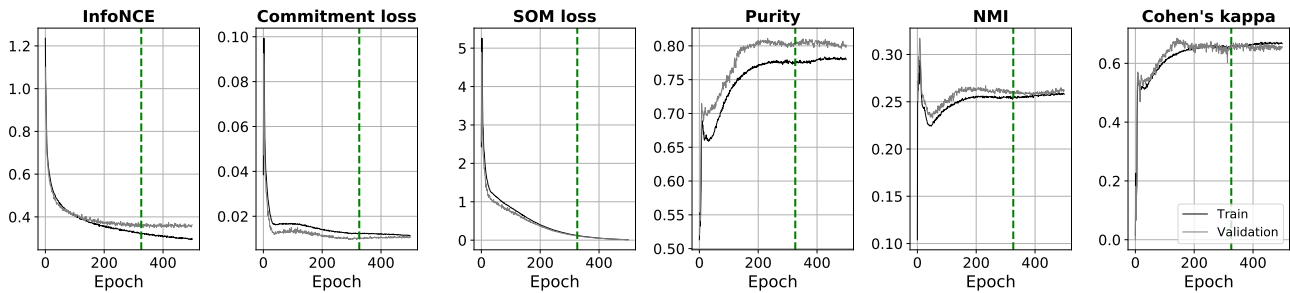


Figure 7: Training curves of SOM-CPC (the model indicated with a * in Table 5) for both the training and validation set. The green dashed line indicates the epoch of the used model, i.e. the one with the lowest validation InfoNCE loss. It can be seen that the epoch with the best clustering and classification performance does not necessarily align with the epoch that has the lowest loss commitment and/or SOM loss.

indicates on average larger jumps over the SOM map through time, which might be caused by the fact that the SimCLR task objective does not incorporate temporal information, while InfoNCE does exploit this. Training time of SOM-SimCLR was, moreover, considerably longer than SOM-CPC with the same settings due to the additional augmentations that need to be computed for every data window and its negative samples.

Figure 7 shows training curves of the training and validation set for the SOM-CPC model that is indicated with a * in Table 5. The green line indicates the epoch with the lowest InfoNCE validation loss. These graphs show that the performance of InfoNCE, the commitment and SOM loss, and classification metrics do not necessarily align, making it dependent on your final goal with the SOM-CPC model what is the most appropriate stopping-criterion.

A.5. Audio experiments

A.5.1. TRAINING DETAILS

Audio streams were encoded in windows of 0.01 seconds (=160 samples). For CPC, GRU-DESOM, and SOM-CPC the contextual information of $L = 127$ previous windows was aggregated using a GRU, equivalent as proposed by Oord et al. (2019). The InfoNCE objective for (SOM-)CPC was computed on top of the last context vector of the GRU. To test different settings for the GRU-desom model, we distinguished GRU-DESOM that decodes only the last window, given the last context vector, and GRU-DESOM that decodes the full sequence of 128 windows from the respective context vectors.

Table 6 provides the model details of the encoder, and the decoder for the DESOM benchmarks. The reconstruction loss of the DESOM model was found to be hampered in its optimization when using the encoder architecture, as adopted for the SOM-CPC model. As such, the downsampling factor of the encoder was reduced for the DESOM model to enable minimization of the task loss during training.

Table 6: Model details for the audio experiments in Section 4.3.

Layer type	Output size	Channels	Activation	Kernel size	Strides	Dilation	Padding
Encoder for SOM-CPC and CPC							
Conv1D	$bs \times 512 \times 32$	512	ReLU	10	5	1	3
Conv1D	$bs \times 512 \times 8$	512	ReLU	8	4	1	2
Conv1D	$bs \times 512 \times 4$	512	ReLU	4	2	1	1
Conv1D	$bs \times 512 \times 2$	512	ReLU	4	2	1	1
Conv1D	$bs \times 512$	512	ReLU	4	2	1	1
GRU	$bs \times 512$	512	-	-	-	-	-
Encoder for (GRU-)DESOM							
Conv1D	$bs \times 512 \times 32$	512	ReLU	10	5	1	3
Conv1D	$bs \times 512 \times 8$	512	ReLU	8	4	1	2
Conv1D	$bs \times 512 \times 4$	512	ReLU	4	2	1	1
Conv1D	$bs \times 512 \times 2$	512	ReLU	4	2	1	1
Conv1D	$bs \times 512 \times 2$	512	ReLU	4	1	1	same
Flatten	$bs \times 1024$	-	-	-	-	-	-
GRU	$bs \times 1024$	1024	-	-	-	-	-
Decoder for (GRU-)DESOM							
Unflatten	$bs \times 512 \times 2$	-	-	-	-	-	-
Conv1D	$bs \times 512 \times 2$	512	ReLU	4	1	1	same
ConvTranspose1D	$bs \times 512 \times 4$	512	ReLU	4	2	1	1
ConvTranspose1D	$bs \times 512 \times 8$	512	ReLU	4	2	1	1
ConvTranspose1D	$bs \times 512 \times 32$	512	ReLU	8	4	1	2
ConvTranspose1D	$bs \times 1 \times 160$	512	ReLU	10	5	1	3 (+ output pad = 1)

For training of SOM-CPC, we followed the settings from Oord et al. (2019) and set $P = 12$. The number of negative samples was set to $N = 10$, which were drawn randomly from the entire training set. All deep-SOM models were trained for maximally 3000 epochs, using the Adam optimizer (Kingma & Ba, 2015) with a learning rate of $1e-4$ and a batch size of 8. One epoch was defined as a push through of one sequence of 128 windows (or 1 window for DESOM) from each recording. The best model was selected based on the lowest validation task loss.

The supervised linear classifier and disjoint SOM training on top of the frozen CPC embeddings were trained with a batch size of 128, and a learning rate of $1e-4$ and $1e-2$, respectively. Both models were stopped upon convergence of the validation loss (which was after 200 and 250 epochs, respectively).

A.5.2. EXTENDED RESULTS

Quantitative results of the audio experiments can be found in Table 7. For this application, Cohen’s kappa is computed as the average over all data windows in the test set, which is different from the synthetic and physiological case, where it was computed as the average and one standard deviation across recordings. In these audio experiments, all windows from one recording contain the same speaker id label. Computing Cohen’s kappa per-recording, i.e. with having the same label for all windows in that recording, is therefore inappropriate as the computation can not correct for correctness by chance. Table 7 show that SOM-CPC clearly outperformed all variants of the (GRU-)DESOM model, and feature extraction using CPC, followed by PCA and linear or non-linear classification.

Ablations with respect to the temperature τ and the similarity metric in the loss function indicated that simply changing the temperature value did hardly affect SOM-CPC performance. However, changing the similarity metric to be the cosine similarity caused a drop in clustering and classification performance, but only when using a temperature value of 1. This result is in line with the experimental findings in the physiological case (see Appendix A.4.3), and suggests that a low temperature value - which was found beneficial in the SimCLR objective (Chen et al., 2020; Woo et al., 2022) - did not have

Table 7: Test set performance of various models trained on audio recordings. SOMs of models with a * are visualized in Figure 5. Bold values indicate the best performance per column (excluding the upper bound of the vanilla CPC model, which does not result in a 2D representation).

Model	α	\mathcal{S}	$\mathcal{L}_{\text{SOM}} \text{sg}[\cdot]$	Purity	NMI	Cohen's kappa	TE	
CPC + linear classifier	-	-	-	-	-	1.00	-	
CPC + K-means	-	-	-	1.00	.60	1.00	-	
CPC ($F = 2$) + linear classifier	-	-	-	-	-	.00	-	
CPC ($F = 2$) + K-means	-	-	-	.13	.01	.03	-	
CPC + PCA + linear classifier	-	-	-	-	-	.86	-	
CPC + PCA + K-means	-	-	-	.89	.54	.88	-	
DESOM	1e-5	Gaussian	\times	.18	.03	.06	.14 \pm .04	
	1e-4	Gaussian	\times	.23	.08	.11	.34 \pm .05	
	1e-3	Gaussian	\times	.31	.13	.20	.68 \pm .05	
	1e-2	Gaussian	\times	.13	.00	.00	1.00 \pm .00	
GRU-DESOM (reconstructing last window)	1e-5	Gaussian	\times	.19	.04	.08	.13\pm.03	
	1e-4	Gaussian	\times	.26	.09	.15	.20 \pm .04	
	1e-3	Gaussian	\times	.31	.13	.21	.42 \pm .08	
	1e-2	Gaussian	\times	.32	.14	.22	.46 \pm .06	
GRU-DESOM (reconstructing full sequence)	1e-5	Gaussian	\times	.19	.05	.08	.34 \pm .05	
	1e-4	Gaussian	\times	.30	.12	.20	.59 \pm .07	
*	1e-3	Gaussian	\times	.33	.14	.22	.78 \pm .05	
	1e-2	Gaussian	\times	.29	.12	.19	.57 \pm .07	
SOM-CPC (ours)	1e-5	Gaussian	\times	.99	.73	.99	.14 \pm .08	
	1e-4	Gaussian	\times	1.00	.63	1.00	.24 \pm .12	
	1e-3	Gaussian	\times	1.00	.61	1.00	.33 \pm .10	
	1e-2	Gaussian	\times	1.00	.61	.99	.33 \pm .10	

Ablations	1e-3	Gaussian	\checkmark	1.00	.61	.99	.28 \pm .08	
	1e-2	Gaussian	\checkmark	1.00	.61	.99	.35 \pm .10	
	.1	Gaussian	\checkmark	1.00	.61	1.00	.35 \pm .10	
	1	Gaussian	\checkmark	1.00	.61	1.00	.38 \pm .11	
	SOM-CPC ($\tau = 0.07$, sim = cosine sim.)	1e-3	Gaussian	\times	.99	.61	.99	.42 \pm .12
	SOM-CPC ($\tau = 1$, sim = cosine sim.)	1e-3	Gaussian	\times	.88	.55	.86	.17 \pm .06
	SOM-CPC ($\tau = 0.07$, sim = dot prod.)	1e-3	Gaussian	\times	1.00	.61	.99	.38 \pm .10
	CPC + SOM (disjoint)	-	Gaussian	-	1.00	.62	1.00	.28 \pm .11




Figure 8: Projecting the test set on the 2D PCA space after CPC (with $F = 128$) encoding, shows no division of the green and the red clusters in two sub-clusters, something that is visible in the SOM of the SOM-CPC model. These sub-clusters were found to relate to recordings that were made with different room acoustics.

much influence on SOM-CPC performance when using the dot product as the similarity metric.

Figure 8 compares the test set projection on the 2D PCA space, created on the CPC features (with $F = 128$), to the SOM from the SOM-CPC model that is denoted with a * in Table 7 (and also visualized in Figure 5-right). This SOM reveals two separate clusters both for speaker 2 (green) and 3 (red), which were found to originate from recordings with different room acoustics (see Section 4.3). This division between recordings of the same speaker is not visible in the 2D PCA projection (Figure 8-left)