# SELF-SUPERVISED DYNAMICAL SYSTEM REPRESEN-TATIONS FOR PHYSIOLOGICAL TIME-SERIES

## **Anonymous authors**

000

001

002003004

006 007 008

009 010

011

012

013

014

015

016

017

018

019

021

024

025

026027028

029

031

033

034

037

040

041

042

043

044

046

047

048

050 051

052

Paper under double-blind review

#### **ABSTRACT**

The effectiveness of self-supervised learning (SSL) for physiological time series depends on the ability of a pretraining objective to preserve information about the underlying physiological state while filtering out unrelated noise. However, existing strategies are limited due to reliance on heuristic principles or poorly constrained generative tasks. To address this limitation, we propose a pretraining framework that exploits the information structure of a dynamical systems generative model across multiple time-series. This framework reveals our key insight that class identity can be efficiently captured by extracting information about the generative variables related to the system parameters shared across similar time series samples, while noise unique to individual samples should be discarded. Building on this insight, we propose PULSE, a cross-reconstruction-based pretraining objective for physiological time series datasets that explicitly extracts system information while discarding non-transferrable sample-specific ones. We establish theory that provides sufficient conditions for the system information to be recovered, and empirically validate it using a synthetic dynamical systems experiment. Furthermore, we apply our method to diverse real-world datasets, demonstrating that PULSE learns representations that can broadly distinguish semantic classes, increase label efficiency, and improve transfer learning.

# 1 Introduction

Self-supervised learning (SSL) is a powerful framework for learning general-purpose representations from unlabeled datasets. These representations may be especially powerful for physiological time-series, making it possible to track physiological states, detect diseases, and improve our understanding of biology (Perochon et al.; Chen et al., 2023b; Li et al., 2025a). The success of an SSL approach here is determined by how well its pretraining objective guides an encoder to preserve information relevant to the identity of the underlying physiological process while filtering out irrelevant noise (Tian et al., 2020).

Although many time-series SSL methods have been proposed, their strategies often make heuristic assumptions that may poorly match the characteristics of physiological signals. For instance, time-series contrastive learning (CL) aims to learn representations that encode information shared between two views of the data while remaining invariant to unshared information. However, positive pairs formed through augmentations (e.g. jittering, scaling) can incorrectly relate samples with different clinical diagnoses, and sampling heuristics (e.g. temporal-proximity) may fail to capture nonstationarity, where physiological states may shift abruptly over time. Masked autoencoding (MAE) takes a different approach, using a reconstruction task to recover information shared between masked and unmasked regions while ignoring unshared features (Kong and Zhang, 2023). Yet, similar to CL, masking strategies (e.g., random, patch, channel, frequency) are heuristically designed without regard to the data-generating process, which can inadvertently treat important temporal relationships as noise, thereby misrepresenting the underlying process. Consequently, SSL methods that rely on heuristics may fail to preserve information essential for distinguishing physiological states which limits their effectiveness when applied to real-world clinical tasks (Xiao et al., 2020).

Sequential variational autoencoders (SVAEs) offer an alternative to heuristic SSL by defining an autoencoding task based on the characteristics of physiological time-series. Specifically, SVAEs define a dynamical systems data-generating process and use structured decoders to guide learned representations toward components of this process. These structural constraints allow SVAEs to

056

060 061

062

063

064

065

066

067

068

069

071

072

073

074

075

076

077

078

079

081

083

084

085

086

087

880

089

090

091

092

094

095

096

098

099

100

101

103 104 105

106

107

Figure 1: Intuition behind PULSE. A dynamical systems model of a physiological time-series dataset allows us to distinguish between information that is shared between similar time series and information that is sample-specific and not transferrable. PULSE leverages this distinction to learn representations that preserve shared system information while discarding sample-specific ones.

capture information about time-series dynamics that heuristic methods may ignore or distort. This makes them broadly effective for physiological time-series, where information about the dynamics drive the signal's evolution and directly reflect the underlying state (Nayak et al., 2018; Sherman, 2011; Marmarelis, 2004; McKenna et al., 1994; Brooks et al., 2021).

However, SVAEs have two critical limitations that prevent them from consistently outperforming heuristic SSL methods. First, the autoencoding task does not separate signal from noise, as both are reconstructed jointly. For physiological time series, this distinction is critical because not all observed information reflects the underlying state. For example, the initial time or amplitude of a recording is much less informative of a signal's identity than its temporal behavior, and transient fluctuations often occur naturally without indicating changes in the physiological process. Second, because samples are reconstructed independently, there is no mechanism to learn information shared across different samples. Since physiological time series often exhibit repeated patterns, capturing this shared information is important because it can relate different samples with similar behaviors.

In this work, we address the limitations of heuristic SSL and dynamical systems models for physiological time-series by introducing PULSE (Physiological self-sUpervised Learning using System Encoders). This is accomplished by introducing a novel pretraining objective that selectively retains information relevant to the underlying process based on structure available from a dynamical systems generative model. In our framework, we conceptualize a time series dataset as the result of a hierarchical graphical model where related samples are generated from the same underlying dynamical system. Under this model, we obtain our key insight as illustrated in Figure 1: system information from the generative parameters should be preserved since it relates samples produced by similar processes, whereas sample-specific information about factors that are unique to each sample, such as initial conditions and process noise, is non-transferrable and should be discarded. To operationalize this idea, we propose a practical cross-reconstruction task that encourages the learned system representation to be invariant to sample-specific factors by reconstructing multiple similar time series across randomly sampled initial conditions. We then present theory on sufficient conditions under which system information is recovered and empirically validate this theory through synthetic experiments. Finally, in several real-world datasets and tasks, we demonstrate that PULSE achieves consistent performance improvements over recent CL and MAE methods. Our contributions are summarized as follows:

- 1. This is the first work to develop an SSL pretraining approach that uses the information structure in dynamical systems to explicitly separate relevant and irrelevant information. We formalize this distinction through a hierarchical generative model over a time-series dataset and propose a practical pretraining strategy that selectively extracts the desired information.
- 2. We introduce theory for PULSE that provides conditions for when system information is recovered and empirically validate this theory on a synthetic dynamical systems experiment.
- 3. In many real-world physiological applications, PULSE achieves SOTA performance, outperforming a representative set of baselines in linear evaluation, data efficiency, and transferability.

#### 2 BACKGROUND AND RELATED WORK

**Self-Supervised Learning for Time Series.** CL and MAE are the most studied SSL paradigms for time-series due to their success in computer vision (Gui et al., 2024; Li et al., 2024; 2025b). In CL, positive pair design is crucial and determines what information the representation is invariant to (Tian et al., 2020). For time series, positive pairs created via augmentations like scaling,

jittering, or adding noise as in SimCLR (Chen et al., 2020) often yield inconsistent performance across datasets (Liu et al., 2024). Positive pairs can also be constructed via sampling methods such as selecting time neighbors in TNC (Tonekaboni et al., 2021), creating two crops of a single segment in TS2vec (Yue et al., 2022), or through a learned reconstruction measure as in REBAR (Xu et al., 2023). In MAE, the masking strategy plays a central role in determining what shared information the model must recover to solve the masked reconstruction task (Kong and Zhang, 2023). TimeMAE (Cheng et al., 2023) introduces block masking to capture shared information across subseries. PatchTST (Nie et al., 2022) applies a similar idea termed patch masking, and further shows that assuming channel independence can improve forecasting performance. Wang et al. (2024) explores various time and channel masking strategies, but finds their performance to be highly dataset-dependent. CiTrus (Geenjaar and Lu, 2025) proposes a MAE task that predicts either spectrogram or time-series, but shows that neither approach is consistently effective across a range of biosignals. Sensitivity to dataset characteristics limits these approaches when the dataset properties are not well understood a priori. Instead, PULSE is based on a dynamical systems model across multiple time series and demonstrates consistent performance in many domains.

**Dynamical Systems Models of Physiological Time Series.** Many physiological time-series arise from underlying physical and biochemical processes that can be well described by dynamical systems models. Thus, dynamical systems provides a shared modeling framework useful for describing a broad range of physiological processes. Since a physiological system evolves independently of how it is measured, the observations  $y_t = g_y(x_t)$  can be viewed as the output of an observation function applied to a latent physiological state  $x_t \in \mathbb{R}^n$ , which itself evolves over time according to the dynamics  $x_{t+1} = g_x(x_t, \Theta)$ , where  $g_x$  is parameterized by system variables  $\Theta$ . This defines a generative process known as the *state-space model* (SSM). In this framework, a time series is generated by specifying an initial condition  $x_{t_0}$ , evolving it forward in time according to  $g_x$ , and producing measurements according to  $g_y$ . Although dynamical systems models have been successfully used to study a wide range of physiological processes, including cardiac dynamics (Bianco et al., 2024), brain activity (Chen et al., 2024; Mudrik et al., 2024), and motor control (Shenoy et al., 2011), they have not yet been used to guide which information should be preserved or discarded in a SSL objective. In this work, we demonstrate that exploiting dynamical system structure in time-series SSL pretraining can lead to consistent improvements in representation learning for physiological signals.

Sequential Variational Autoencoders. SVAEs can be viewed as a time-series SSL strategy that leverages the state-space formulation to learn representations corresponding to components of a dynamical system models. Training is performed through an autoencoding task by maximizing the ELBO,  $\mathbb{E}_{q(z|y)}[\log p(y|z)] - \beta \mathrm{KL}(q(z|y)||p(z))$ , where z are latent variables, q(z|y) is the approximate posterior, p(y|z) is the likelihood, p(z) is the latent prior, and  $\beta$  is the regularization strength for the KL term. The encoder q(z|y) and decoder p(y|z) are often parameterized by neural networks, and are structured to reflect specific generative assumptions. For instance, LFADS (Sussillo et al., 2016) encodes each time series into an initial condition  $x_{t_0}$  that evolves under a shared dynamics function and uses a GRU decoder to reconstruct the observed data. DSVAE (Li and Mandt, 2018) assumes that the dynamics is data-dependent and can be factorized into static and dynamic generative factors. Unlike CL and MAE, which aims to selectively recover useful sources of information during pretraining, SVAE methods share the common limitation that the  $\mathbb{E}_{q(z|y)}[\log p(y|z)]$  term does not distinguish between meaningful signal from noise and instead encourages explaining all observed variability in the data. Consequently, the learned representations can be overly sensitive to irrelevant information, reducing their transferability to downstream tasks.

## 3 PULSE APPROACH

In this section, we present PULSE. First, we develop a dynamical systems generative model for multiple time-series samples to identify which information is transferrable and which is not. Then, we introduce a practical pretraining strategy to extract the transferrable information. Finally, we provide theoretical guarantees for conditions under which the transferrable information is recovered.

#### 3.1 REVEALING SHARED INFORMATION BETWEEN MULTIPLE TIME-SERIES

To determine which information should be retained or discarded, we first need to characterize the information present in a physiological time-series dataset. We do this by introducing notation to describe the data and constructing a graphical model of its generative process, which makes explicit the sources of information available in the observations.

**Notation.** A physiological time-series dataset often consists of long continuous time series recordings that are segmented into shorter, fixed-length windows. Let  $\mathbf{D} \in \mathbb{R}^{R \times T \times M}$  denote a dataset of R recordings, each consisting of T time steps and M measurement channels. We use  $\mathbf{D}_{r,\tau:h}$  to denote a subsequence from the rth recording, starting at time index  $\tau$  and ending at index h. We can then construct a dataset of N time-series windows,  $\mathbf{Y} \in \mathbb{R}^{N \times W \times M}$ , where each window is defined as  $\mathbf{Y}_i = \mathbf{D}_{r_i,\tau_i:\tau_i+W}$  for a window size W, with  $r_i$  and  $\tau_i$  denoting the recording and start index of the ith sample, respectively. A specific time-slice is denoted as  $\mathbf{Y}_{i,t_1}$ , where  $t_1 \in \{1,\ldots,W\}$ , and consecutive time-slices are defined as  $t_{k+1} = t_k + 1$ , where  $t_k = t_k + 1$ .

Time-Series Dataset Generative Model. As shown in Figure 2, we model Y with a dynamical systems generative model, where each  $Y_i$ is the output of an SSM with system parameters  $\Theta_i$ , and an initial condition  $\mathbf{X}_{i,t_0}$ . While each  $Y_i$  could in principle be generated by a unique system, physiological activity is often stereotyped, with many underlying processes exhibiting consistent, repeatable patterns over time. For example, a walk cycle captured by an accelerometer displays repeated phases such as heel strike, mid-stance, and toe-off, while an ECG signal shows recurring PQRST complexes during normal sinus rhythm. As a result, different samples may share the same underlying system, and the number of unique  $\Theta_i$  is generally smaller than N. To capture this, we modify

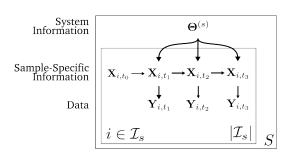


Figure 2: Our graphical model of multiple timeseries windows, based on dynamical systems, distinguishes transferable system information shared across similar time-series from non-transferable information unique to each sample such as initial conditions and process noise.

the notation such that  $\Theta^{(s)}$  represents the system parameters corresponding to all indices in  $\mathcal{I}_s$ , the set of samples generated by system  $s \in \{1, \ldots, S\}$ . In other words,  $\Theta_i = \Theta^{(s)}$  for all  $i \in \mathcal{I}_s$ . Thus, the joint distribution for this generative model is given by,

$$p(\mathbf{Y}, \mathbf{X}, \mathbf{\Theta}) = \prod_{s=1}^{S} \prod_{i \in \mathcal{I}_s} p(\mathbf{X}_{i, t_0}, \mathbf{\Theta}^{(s)}) \left[ \prod_{k=1}^{W} p(\mathbf{Y}_{i, t_k} | \mathbf{X}_{i, t_k}) \right] \left[ \prod_{k=2}^{W} p(\mathbf{X}_{i, t_k} | \mathbf{X}_{i, t_{k-1}}, \mathbf{\Theta}^{(s)}) \right], \quad (1)$$

where  $p(\mathbf{X}_{i,t_0}, \mathbf{\Theta}^{(s)}) = p(\mathbf{X}_{i,t_0})p(\mathbf{\Theta}^{(s)})$  (i.e.  $\mathbf{X}_{i,t_0}$  and  $\mathbf{\Theta}^{(s)}$  are independent). The observation and transition densities follow the SSM, with  $p(\mathbf{Y}_{i,t_k}|\mathbf{X}_{i,t_k})$  given by  $\mathbf{Y}_{i,t_k} = g_y(\mathbf{X}_{i,t_k}) + \epsilon_{i,t_k}$  and  $p(\mathbf{X}_{i,t_k}|\mathbf{X}_{i,t_{k-1}},\mathbf{\Theta}^{(s)})$  given by  $\mathbf{X}_{i,t_k} = g_x(\mathbf{X}_{i,t_{k-1}},\mathbf{\Theta}^{(s)}) + \nu_{i,t_k}$ , where  $\epsilon$  and  $\nu$  are noise terms.

The factorization in Eq. 1 reveals a hierarchy of information in Y. One source of information comes from the system variables  $\Theta^{(s)}$ , which govern the evolution of the time series and are shared across  $Y_i$  with the same underlying dynamics, making it transferable between different samples. Learning a representation that captures this system information produces a space in which samples with similar dynamics are naturally grouped together. Another source of information is unique to each  $Y_i$ , such as the initial value  $X_{i,t_0}$ , as well as observation and dynamics noise  $\epsilon$  and  $\nu$ . In our setting, a representation should be invariant to this information, as it is not shared across  $Y_i$  and cannot be transferred between different samples. This insight is especially relevant for physiological time series, where a signal's identity is determined by the underlying system dynamics rather than by the exact starting value, sensor noise, or transient fluctuations. Therefore, by extracting system information and discarding sample-specific information, we obtain a representation that can relate physiological time series based on their temporal characteristics while ignoring irrelevant factors.

# 3.2 PULSE PRETRAINING FOR RECOVERING SYSTEM INFORMATION

Our goal is to design a practical pretraining strategy that encourages an encoder to recover system information while ignoring sample-specific factors. A promising strategy is a dynamical systems cross-reconstruction task, where given two samples from the same system (i.e.,  $\mathbf{Y}_i$  and  $\mathbf{Y}_j$  with  $i,j \in \mathcal{I}_s$ ), the system information inferred from  $\mathbf{Y}_i$  is used to reconstruct an independently realized sample  $\mathbf{Y}_j$ . By requiring the system representation from  $\mathbf{Y}_i$  to reconstruct multiple random samples of  $\mathbf{Y}_j$ , an encoder is encouraged to keep only the shared information between these samples. According to Eq. 1, the only shared variables are  $\mathbf{\Theta}^{(s)}$ , and encoding irrelevant factors may lead to poor reconstruction.

Cross-Reconstruction with Similar Pairs. To formalize this approach, we define an inference step and generative process given sample pairs  $(\mathbf{Y}_i, \mathbf{Y}_j)$  produced by  $\mathbf{\Theta}^{(s)}$  as input. For inference, we introduce two encoders that separate transferable from non-transferable dynamical systems components: a system encoder  $f_{\rm sys}$  to estimate shared system information and an initial condition encoder  $f_{\rm init}$  to estimate the sample-specific initial condition. As shown in Figure 3,  $f_{\rm sys}$  uses the dilated convolution (Yue et al., 2022) to extract information across the entire window according to  $\mathbf{\Theta}_i = f_{\rm sys}(\mathbf{Y}_i)$ . In contrast,  $f_{\rm init}$  is implemented as a 2-layer CNN whose receptive field is centered around a specified time  $t_0$ , producing  $\mathbf{X}_{j,t_0} = [f_{\rm init}(\mathbf{Y}_j)]_{t_0}$  where  $t_0 = 1$  selects the initial condition needed to reconstruct a sample from the first time step until the end of the window. For generation, we use an SSM decoder and define the cross-reconstruction objective,

$$\mathcal{L}_{\text{Cross}}(\mathbf{Y}_i, \mathbf{Y}_j) = \mathbb{E}_{(\mathbf{Y}_i, \mathbf{Y}_j) \sim \mathcal{P}} \left[ \sum_{k=1}^{W} \|\mathbf{Y}_{j, t_k} - g_y(g_x(\mathbf{X}_{j, t_{k-1}}, \mathbf{\Theta}_{i, t_k}))\|^2 \right], \tag{2}$$

where  $\mathcal{P}$  is a distribution over sample pairs  $i,j\in\mathcal{I}_s$ , and  $g_x$  and  $g_y$  are parameterized by a GRU and linear projection layer respectively. Here, we implement  $\Theta_i$  as the GRU input, since its hidden state evolves according to dynamics defined by input-dependent gating. Additionally, following prior dynamical systems methods (Li and Mandt, 2018), we further factorize  $\Theta_i$  into separate time-invariant and time-varying components to represent nonstationary behaviors, a key feature of physiological time-series. To do this, we decompose  $\Theta_i$  into a time-invariant component  $\theta_i$ , obtained via max pooling over time, and a time-varying component  $\tilde{\theta}_{i,t_k}$ , obtained via a two-layer CNN over the latent dimension, and then concatenate the result at each time step to form  $\Theta_{i,t_k} = [\theta_i, \tilde{\theta}_{i,t_k}]$ .

To minimize Eq. 2,  $f_{\rm sys}$  must extract shared information from  $\mathbf{Y}_i$  that can explain the evolution of  $\mathbf{Y}_j$ . This means encoding only the underlying system variables  $\mathbf{\Theta}^{(s)}$ , since these are the factors that remain invariant across samples from the same system. Furthermore, because  $\mathcal{P}$  involves optimizing over randomly chosen pairs,  $f_{\rm sys}$  cannot rely on sample-specific information from  $\mathbf{Y}_i$ , since this information will not be present in  $\mathbf{Y}_j$  and may increase  $\mathcal{L}_{\rm Cross}$  if present. Thus, the  $\mathcal{L}_{\rm Cross}$  loss encourages the encoder to discard non-shared factors and focus only on shared system information.

Cross-Reconstruction with PULSE Pseudo-Pairs. Unfortunately, a problem in Eq. 2 is that, in an unlabeled dataset, we do not have access to  $\mathcal{I}_s$  and therefore cannot sample from  $\mathcal{P}$ . Instead, we propose a sampling strategy that constructs semantically similar time series pseudo-pairs  $(\mathbf{Y}_i, \mathbf{Y}_i)$  from a single sample. Since  $\Theta^{(s)}$  are independent of the initial condition in Eq. 1 and should define dynamics that can generate time series from any starting state, representations that capture system information should generalize across multiple initial conditions. However, initial conditions cannot be chosen arbitrarily, as data may not be available to supervise reconstruction from arbitrary starting points. To address this, we select an initial condition at a random time step

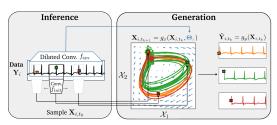


Figure 3: PULSE aims to recover system information through an inference process that uses two encoders,  $f_{\rm sys}$  to estimate shared parameters of a latent dynamical systems and  $f_{\rm init}$  to estimate sample-specific initial conditions. By requiring  $\Theta_i$  to support reconstruction of randomly sampled  $\mathbf{X}_{i,t_0}$ , we encourage the recovered system information to be invariant to the sample-specific ones.

 $t_0 \sim \text{Uniform}(1, W/2)$  within each window  $\mathbf{Y}_i$ , and task the model with reconstructing the segment from  $t_0$  to the end. Note we restrict  $t_0$  to the first half of the window to prevent overfitting to short subsequences. Thus, for any window  $\mathbf{Y}_i$ , we generate a semantically similar time series  $\widetilde{\mathbf{Y}}_i = \mathbf{Y}_{i,t_0:W}$  to serve both as a pair for estimating sample-specific information and as a cross-reconstruction target. By requiring  $\boldsymbol{\Theta}_i$  to support accurate reconstruction across varying initial conditions, the encoder learns to recover only system information while remaining invariant to information about the initial condition. This leads to the PULSE pretraining objective,

$$\mathcal{L}_{\text{PULSE}}(\mathbf{Y}_i) = \mathbb{E}_{t_0} \left[ \sum_{k=1}^{W} \| \widetilde{\mathbf{Y}}_{i,t_k} - g_y(g_x(\mathbf{X}_{i,t_{k-1}}, \boldsymbol{\Theta}_{i,t_k}) \|^2 \right], \tag{3}$$

where  $\mathbf{X}_{i,t_0} = [f_{\text{init}}(\mathbf{Y}_i)]_{t_0}$  and  $\mathbf{\Theta}_i = f_{\text{sys}}(\mathbf{Y}_i)$ . We find that using multiple  $\widetilde{\mathbf{Y}}_i$ 's to estimate the expectation can improve performance, and in our experiments we use up to four samples.

#### 3.3 Provable Recovery of System Information

We now provide a theoretical analysis of our framework and identify conditions under which cross-reconstruction provably recovers system information. Our strategy is to build on prior work showing that MAE pretraining implicitly recovers information from the minimal set of latent variables shared  $^{1}$   $^{1}$  between masked and unmasked regions in a hierarchical data-generating process (Kong and Zhang, 2023). By viewing cross-reconstruction as an MAE task under a specific masking strategy, we can extend this theory to characterize how different masking choices determine the type of information recovered in our own time-series generative model.

Cross-reconstruction can be viewed as an MAE task by treating the pair  $(\mathbf{Y}_i, \mathbf{Y}_j)$  as a single joint input with a pair of masking variables  $(\mathbf{m}_i, \mathbf{m}_j)$ , where each  $\mathbf{m}_i \in \{0,1\}^{W \times M}$  indicates for every element (w,m) whether it is observed  $(m_{i,w,m}=1)$  or masked  $(m_{i,w,m}=0)$ . In this view, Eq. 2 corresponds to setting  $\mathbf{m}_{i,1:W,1:M}=1$  to retain  $\mathbf{Y}_i$  as input, while fully masking the other sample  $\mathbf{m}_{j,1:W,1:M}=0$ , so that  $\mathbf{Y}_j$  is removed and can serve as the reconstruction target. The effect of this masking strategy on the type of information recovered can then be characterized given our generative model in Eq. 1. To make this precise, we outline assumptions on the generative process.

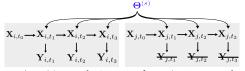
**Assumption 1** (Data-Generating Process). The process in Equation (1) satisfies the following conditions: (i) the fully factorized generative model is a DAG; and (ii) each function  $q_k$  is invertible.

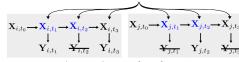
**Assumption Interpretation.** Part (i) ensures that our theory applies to complex systems with elaborate parameter factorizations, as long as they remain acyclic. Part (ii) guarantees that no information is lost during the generative process and is adopted from prior work on identifiable deep generative models (Locatello et al., 2020; Von Kügelgen et al., 2021).

Given this data-generating process and the cross-reconstruction masking scheme described above, we present our theory, which identify the  $\mathcal{C}$  between masked and unmasked regions. This  $\mathcal{C}$  corresponds to the information that is implicitly recovered during MAE pretraining (Kong and Zhang, 2023).

**Theorem 1.** Given two time series  $\mathbf{Y}_i$  and  $\mathbf{Y}_j$  independently sampled from the same system (i.e.,  $\mathbf{\Theta}_i = \mathbf{\Theta}_j = \mathbf{\Theta}^{(s)}$ ) under the generative process defined by Eq. 1 and Assumption 1, the minimal set of latent variables shared is the system parameters  $\mathbf{\Theta}^{(s)}$  if and only if all observables from one series is fully masked (i.e.,  $\mathbf{m}_{i,1:W,1:M} = 0$  and  $\mathbf{m}_{j,1:W,1:M} = 1$ ).

A proof is provided in Appendix A. Theorem 1 states that system information is recovered during a masked reconstruction task when an entire time series is removed from the input pair, as this masking scheme uniquely ensures that the  $\mathcal C$  connecting the masked and unmasked regions contains the system parameters. Importantly, this reconstruction task with whole-sample masking corresponds exactly to the pretraining objective  $\mathcal L_{\text{Cross}}$ . Moreover, this theory predicts that when a time series contains both masked and unmasked regions,  $\mathcal C$  necessarily includes the state variables  $\mathbf X$ , causing the recovered information to confound sample-specific and system information.





A) Positive Mask: System Information Recovered

B) Negative Mask: Sample-Specific Information Recovered

Figure 4: We illustrate how these different masking strategies recover different sources of information in our data-generating process for sample pairs  $(\mathbf{Y}_i, \mathbf{Y}_j)$  where  $i, j \in \mathcal{I}_s$ .  $\mathbf{Y}$  marks an observable that is removed from the input and used as a reconstruction target. Blue highlights  $\mathcal{C}$ , representing the information that is recovered during pretraining. Theorem 1 predicts that  $\mathcal{C} = \{\mathbf{\Theta}^{(s)}\}$  only when information from one sample is fully removed. Gray boxes group latent variables that are specific to each time-series sample.

Since  $\mathcal{L}_{\mathrm{PULSE}}$  can be viewed as an approximation of  $\mathcal{L}_{\mathrm{Cross}}$  that uses pseudo-pairs to simulate independent samples, this theory offers an explanation for how  $\mathcal{L}_{\mathrm{PULSE}}$  recovers system information.

<sup>&</sup>lt;sup>1</sup>Our minimal set of shared latent variables  $\mathcal{C}$  is defined in Theorem 1 of (Kong and Zhang, 2023) as c.

# 4 SYNTHETIC DYNAMICAL SYSTEMS EXPERIMENTS

Set up. We investigate how our findings in Theorem 1 extrapolate to settings where Assumption 1 may not completely hold. Specifically, we consider a synthetic dynamical systems experiment, where  $g_x$  can not be practically inverted due to system chaos. The goal of this task is to learn a representation space that can distinguish between parameter settings while remaining robust to increasing levels of dynamical system noise. This task is motivated by real-world scenarios in which parameter changes in the underlying system may correspond to shifts in physiological state, and effective representations must reliably capture these shifts. Time series data are generated from three stochastic differential equations (Lorenz, Thomas, and Hindmarsh-Rose) across a grid of parameters in bifurcation regions and noise levels  $\sigma = \{0, 1, 3, 5\}$ . Datasets are then constructed by randomly selecting five parameter settings to form a five-class classification problem, with trials split into 70:15:15 for train, validation, and test sets and subsequently segmented into windows where W = 100. For each system at each noise level, we average classification accuracy over 10 random seeds over both dataset sampling and model initialization. The reported values are averaged across all three systems. Full dataset details are provided in Appendix B.

**Results.** Table 1 shows that PULSE consistently achieves the highest classification accuracy among all practical baselines (described in Section 5 and Appendix C), even as  $\sigma$  increases. This demonstrates that PULSE pretraining on pseudo-pairs  $(\mathbf{Y}_i, \widetilde{\mathbf{Y}}_i)$  is more robust to dynamical system noise and can extract class-discriminative features that reveal changes in system parameters under noisy conditions. Furthermore, we validate Theorem 1 by considering a PULSE Oracle where label information is used to identify true  $(\mathbf{Y}_i, \mathbf{Y}_j)$  pairs to construct the positive and negative set ups illustrated in Figure 4. Specifically, the positive model leverages labels to select pairs in  $\mathcal{L}_{\text{Cross}}$ , whereas the negative model applies random temporal masking to each pair before inputting both into  $f_{\text{init}}$  and  $f_{\text{sys}}$ . According to Theorem 1, the negative oracle captures sample-specific information that is uninformative for classifying system parameters, so the positive oracle is expected to consistently outperform it.

σ								PULSE Oracle		
0	SimCLR	TS2Vec	REBAR	PatchTST	TimeMAE	LFADS	DSVAE	PULSE	Positive	Negative
0	93.08	98.68	98.90	77.59	96.29	99.06	98.56	99.58	99.29	98.86
1	83.10	93.07	93.36	50.36	93.42	93.02	90.84	96.09	97.26	96.66
3	70.05	79.78	79.36	39.88	75.08	79.03	76.70	83.42	89.00	84.62
5	62.29	73.67	72.37	37.82	66.63	71.33	69.90	77.34	82.65	76.90

Table 1: Our results confirm Theorem 1's predictions: The positive oracle improves classification accuracy over PULSE pretraining without labels, while the negative oracle reduces performance relative to the positive model and, at  $\sigma=5$ , performance even falls below the practical PULSE algorithm. Moreover, PULSE is the most effective practical algorithm, being the best at distinguishing parameter changes across all  $\sigma$ . **Black bold** indicates the best practical algorithm (pretrained without labels), while **blue bold** indicates the best oracle (pretrained with labels) when it exceeds practical methods. Fig. 4 illustrates the masking strategy used in the positive and negative oracle.

# 5 REAL PHYSIOLOGICAL DATA EXPERIMENTS

In this section, we describe our experimental set up for comparing PULSE pretraining with other SSL pretraining approaches across several physiological datasets and downstream tasks.

**Data.** We consider 4 commonly used physiological time-series datasets from distinct sensor domains, each consisting of long trials with time-varying classification labels. We use Human Activity Recognition (HAR) (Reyes-Ortiz et al., 2015), where human activity is estimated from accelerometer and gyroscope signals; PPG (Schmidt et al., 2018), where optical blood volume signals are used to estimate stress levels; ECG (Moody, 1983), where the heart's electrical activity is used to detect rhythm abnormalities; and EEG (Kemp et al., 2000), where the brain's electrical activity is used to estimate sleep stages. A detailed description of these datasets is provided in Appendix D.

**Baselines.** We benchmark performance against a representative set of SSL pretraining approaches, including three CL methods: SimCLR (Chen et al., 2020), TS2Vec (Yue et al., 2022), and RE-BAR (Xu et al., 2023), two SVAE models: LFADS (Sedler and Pandarinath, 2023; Sussillo et al., 2016) and DSVAE (Yingzhen and Mandt, 2018), and two masked modeling approaches: TimeMAE

	Metric	SimCLR	TS2Vec	REBAR	PatchTST	TimeMAE	LFADS	DSVAE	PULSE
HAR	Accuracy ↑	94.65	93.24	95.35	83.04	92.25	93.55	93.55	93.27
	AUROC ↑	99.38	99.31	99.65	97.44	99.14	99.49	99.36	99.42
_H_	AUPRC ↑	97.63	97.66	98.91	89.90	97.05	98.29	97.69	98.10
r <b>5</b>	Accuracy ↑	34.48	40.23	41.38	59.78	61.35	52.81	58.65	64.27
PPG	AUROC ↑	61.19	64.28	69.77	71.08	78.08	71.10	76.78	80.29
	AUPRC ↑	36.08	39.59	44.57	52.91	56.74	49.59	55.38	59.89
rh	Accuracy ↑	69.92	76.12	81.54	64.40	69.80	61.84	70.42	87.41
ECG	AUROC ↑	82.54	86.56	91.46	70.96	76.61	71.69	82.88	94.93
Ĭ	AUPRC ↑	80.63	85.16	89.85	68.16	76.62	69.21	81.31	94.75
EEG	Accuracy ↑	66.38	83.76	83.71	80.62	83.83	82.43	84.25	85.56
	AUROC ↑	85.45	94.99	95.08	93.55	95.09	94.49	95.42	96.17
	AUPRC ↑	50.95	70.22	70.77	67.37	73.22	68.55	72.25	73.82

Table 2: Linear Probe Classification Results. PULSE performs competitively on HAR and achieves the best results on PPG, ECG, and EEG. Note that while HAR scores are lower than SOTA in this experiment, this representation leads to improved performance in Tables 3 and 4.

(Cheng et al., 2023) and PatchTST (Nie et al., 2022). To evaluate the pretraining objective, we use the same dilated convolution encoder architecture (Yue et al., 2022) across all CL and SVAE experiments, ensuring that performance differences reflect the quality of the pretraining rather than differences in architectures. For MAE baselines, we retain the original transformer encoders, since the architecture is often a critical component of the method and changing it can reduce performance. Additional details on baseline design are provided in Appendix C.

#### 5.1 Linear Probe Evaluation

To assess the ability of pretraining to learn class-discriminative features, we train a linear probe (logistic regression) on the frozen embeddings from each pretrained model to predict the ground truth physiological class labels from each dataset. Table 2 shows that PULSE pretraining achieves strong linear probe performance across all four datasets, performing competitively on HAR and achieving the highest overall scores on PPG, ECG, and EEG. Notably, PULSE substantially outperforms LFADS and DSVAE on ECG and PPG, highlighting that explicitly removing noise provides an clear advantage over standard VAE objectives that do not distinguish between noise and signal. Moreover, the improvements over SimCLR, TS2Vec, and REBAR show that CL's sensitivity to false positive pairs may limit its effectiveness across diverse sensor domains. Finally, the performance increase over PatchTST and TimeMAE demonstrates that designing a generative task to explicitly extract system information produces representations that better distinguish physiological states than those learned by standard masked modeling, which does not leverage this structure.

## 5.2 Semi Supervised Evaluation

Next, we evaluate the label efficiency of the pretrained representations using a semi-supervised classification task on the pretrained frozen embeddings. For each pretrained model, we train a linear probe on 1% and 5% of the ground truth labels and apply Laplace smoothing so that all downstream classes are represented. Reported accuracies are averaged over five random label subsets for each of

Dataset		Supervised	SimCLR	TS2Vec	REBAR	PatchTST	TimeMAE	LFADS	DSVAE	PULSE
	HAR	78.39	71.74	80.57	81.10	33.27	80.79	80.97	79.94	84.74
1 0/	ECG	45.46	63.83	62.77	67.56	57.68	65.15	57.34	67.60	84.77
1 %	PPG	34.38	30.34	31.98	32.33	41.74	40.45	40.22	41.28	42.97
	EEG	77.76	56.72	77.39	77.19	63.45	70.55	74.19	78.40	80.69
	HAR	92.34	85.01	91.76	91.04	54.79	91.55	91.48	90.72	93.14
5 %	ECG	69.20	65.00	63.97	70.12	60.73	68.73	60.04	67.66	84.23
3 %	PPG	42.47	30.62	33.13	38.25	49.48	49.53	45.35	51.15	53.39
	EEG	84.90	61.98	77.09	76.75	73.48	77.50	75.19	78.17	80.45

Table 3: Semi-supervised classification accuracy for 1% and 5% of labels averaged over 25 random seeds. Higher score is better. PULSE outperforms all SSL baselines and most supervised baselines.

the five model initializations from Section 5.1, for a total of 25 seeds. We also include a supervised baseline to estimate performance achievable without pretraining. Table 3 shows that PULSE pretraining consistently outperforms the baselines across all four sensor domains, achieving the highest scores among SSL methods across all datasets. Interestingly, PULSE's HAR representation is more label-efficient, achieving strong performance with very few labels despite lower linear probe scores on the full labeled dataset, suggesting that it captures key class-discriminative features more efficiently. PULSE also outperforms most supervised baselines, highlighting the advantage of its pretrained representations in limited data scenarios.

#### 5.3 Transfer Learning Evaluation

We investigate in-domain transfer learning for classification in two scenarios from Zhang et al. (2022). This task is motivated by real-world applications where we want to transfer knowledge between datasets collected from similar sensors. In the first scenario, a model is pretrained on EEG (Kemp et al., 2000) and fine-tuned on the Epilepsy dataset (Andrzejak et al., 2001). In the second, a model is pretrained on HAR (Reyes-Ortiz et al., 2015) and fine-tuned on Gesture (Liu et al., 2009). Our setup follows Zhang et al. (2022), where we attach a 2-layer MLP head and fine-tune for 40 epochs using the

	EEG -	→ Epilepsy	$HAR \rightarrow Gesture$		
	ACC	AUROC	ACC	AUROC	
SimCLR	93.52	97.52	78.83	93.80	
TS2Vec	93.95	95.87	77.67	95.45	
REBAR	95.27	98.33	78.17	95.54	
PatchTST	95.03	98.04	77.00	95.21	
LFADS	94.71	98.01	78.50	95.42	
DSVAE	94.97	98.17	78.00	95.34	
PULSE	95.82	98.51	83.67	97.20	

Table 4: In-Domain Transfer Learning.

Adam optimizer with a learning rate of 0.0003. A detailed description of the fine-tuning datasets is provided in Appendix D. Table 4 reports accuracy and AUROC averaged over five random seeds for both model initialization and fine-tuning. In both scenarios, PULSE consistently achieves the highest transfer performance, with a substantial gain on the HAR-to-Gesture task. This demonstrates that pretraining designed to explicitly prioritize system information can produce features that transfer effectively to related downstream tasks.

#### 5.4 ABLATION STUDY

Table 5 shows the effect of ablating various PULSE components on the linear probe accuracy across all datasets, reported as  $\Delta = \mathrm{Acc}_{\Delta\mathrm{blated}} - \mathrm{Acc}_{\mathrm{PULSE}}.$  In w/o TV-Params  $\theta_{i,t_k}$ , we remove time-varying parameters and retain only time-invariant ones, i.e.,  $\Theta_i = \theta_i.$  This has the largest effect, resulting in an average accuracy drop of 7.6%, underscoring the importance of modeling non-stationary dynamics in physiological time series. In w/o Sample  $t_0$ , we no longer

	HAR	ECG	PPG	EEG	Avg.
PULSE	93.27	87.41	64.27	85.56	-
w/o TV-Params $\tilde{\theta}_{i,t_k}$ w/o Sample $t_0$	-1.73	-6.83 -7.73	-9.22 -15.96	-12.7 -0.81	-7.62 -6.42
Shared $f_{\text{sys}}$ and $f_{\text{init}}$		-1.75	-1.58	-0.54	-1.22

Table 5: Relative effect of ablation on linear probe accuracy. Baseline accuracy (top row) and relative effect ( $\Delta$ ) of removing components (other rows). Avg. reports the mean  $\Delta$  across datasets.

construct pseudo-pairs by randomly sampling  $t_0$ . Instead, we fix  $t_0 = 1$  during initial condition inference for every sample. This leads to a 6.4% drop in average accuracy, showing that training the system representation to reconstruct multiple related samples is important and that our pseudo-pair strategy effectively constructs relevant samples without labels. In Shared  $f_{\rm sys}$  and  $f_{\rm init}$ , we estimate  $\mathbf{X}_{i,t_k}$  from the output of  $f_{\rm sys}$  such that  $\mathbf{X}_{i,t_k} = [f_{\rm init}(f_{\rm sys}(\mathbf{Y}_i))]_{t_k}$ , rather than using two separate encoders. This results in a 1.22% drop, suggesting that explicitly separating transferrable and non-transferrable information can improve time-series representation learning.

# 6 Conclusion

In this work, we introduced PULSE, an improved SSL pretraining framework that improves performance across diverse tasks by preserving system information and filtering out sample-specific noise. By explicitly using dynamical systems to guide the design of time-series SSL objectives, we hope to inspire future research on formulating SSL pretraining strategies based on an underlying generative model to obtain improved representations of more complex physiological phenomena.

### 7 REPRODUCIBILITY STATEMENT

Upon acceptance, we will release our GitHub code publicly. This will include code for downloading and preprocessing the datasets used in our work, training our approach, and evaluating results. It will also provide baseline implementations and the configuration files used to run our experiments as well as model checkpoints.

# REFERENCES

- Ralph G Andrzejak, Klaus Lehnertz, Florian Mormann, Christoph Rieke, Peter David, and Christian E Elger. Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Physical Review E*, 64(6):061907, 2001.
- M Bianco, A Scarciglia, C Bonanno, and G Valenza. Heart rate variability series is the output of a non-chaotic system driven by dynamical noise. *arXiv preprint arXiv:2404.11385*, 2024.
- George E Billman. Heart rate variability–a historical perspective. *Frontiers in physiology*, 2:86, 2011.
- Justin Brooks, Joshua C Crone, and Derek P Spangler. A physiological and dynamical systems model of stress. *International Journal of Psychophysiology*, 166:83–91, 2021.
- Stanislas Chambon, Mathieu N Galtier, Pierrick J Arnal, Gilles Wainrib, and Alexandre Gramfort. A deep learning architecture for temporal sleep stage classification using multivariate and multimodal time series. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(4): 758–769, 2018.
- Danfeng Chen, Junsheng Li, Wei Zeng, and Jun He. Topology identification and dynamical pattern recognition for hindmarsh–rose neuron model via deterministic learning. *Cognitive Neurodynamics*, 17(1):203–220, 2023a.
- Li-Chin Chen, Kuo-Hsuan Hung, Yi-Ju Tseng, Hsin-Yao Wang, Tse-Min Lu, Wei-Chieh Huang, and Yu Tsao. Self-supervised learning-based general laboratory progress pretrained model for cardiovascular event detection. *IEEE journal of translational engineering in health and medicine*, 12:43–55, 2023b.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmLR, 2020.
- Yenho Chen, Noga Mudrik, Kyle A Johnsen, Sankaraleengam Alagapan, Adam S Charles, and Christopher Rozell. Probabilistic decomposed linear dynamical systems for robust discovery of latent neural dynamics. *Advances in Neural Information Processing Systems*, 37:104443–104470, 2024.
- Mingyue Cheng, Qi Liu, Zhiding Liu, Hao Zhang, Rujiao Zhang, and Enhong Chen. Timemae: Self-supervised representations of time series with decoupled masked autoencoders. *arXiv preprint arXiv:2303.00320*, 2023.
- Mukeshwar Dhamala, Viktor K Jirsa, and Mingzhou Ding. Transitions to synchrony in coupled bursting neurons. *Physical Review Letters*, 92(2):028101, 2004.
- Eloy Geenjaar and Lie Lu. Citrus: Squeezing extra performance out of low-data bio-signal transfer learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 16781–16789, 2025.
- Emile Franc Doungmo Goufo, Melusi Khumalo, and Patrick M Tchepmo Djomegni. Perturbations of hindmarsh-rose neuron dynamics by fractional operators: Bifurcation, firing and chaotic bursts. *Discrete and Continuous Dynamical Systems-S*, 13(3):663–682, 2020.

- Alexandre Gramfort, Martin Luessi, Eric Larson, Denis A. Engemann, Daniel Strohmeier, Christian Brodbeck, Roman Goj, Mainak Jas, Teon Brooks, Lauri Parkkonen, and Matti S. Hämäläinen. MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7(267):1–13, 2013. doi: 10.3389/fnins.2013.00267.
  - Jie Gui, Tuo Chen, Jing Zhang, Qiong Cao, Zhenan Sun, Hao Luo, and Dacheng Tao. A survey on self-supervised learning: Algorithms, applications, and future trends. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9052–9071, 2024.
  - Seongsil Heo, Sunyoung Kwon, and Jaekoo Lee. Stress detection with single ppg sensor by orchestrating multiple denoising and peak-detecting methods. *IEEE Access*, 9:47777–47785, 2021.
  - James L Hindmarsh and RM Rose. A model of neuronal bursting using three coupled first order differential equations. *Proceedings of the Royal society of London. Series B. Biological sciences*, 221(1222):87–102, 1984.
  - Shunsuke Kamiya, Jun Kitazono, and Masafumi Oizumi. Koopman operator based dynamical similarity analysis for data-driven quantification of distance between dynamics. In *ICLR 2024 Workshop on Representational Alignment*, 2024.
  - Bob Kemp, Aeilko H Zwinderman, Bert Tuk, Hilbert AC Kamphuisen, and Josefien JL Oberye. 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.
  - Patrick Kidger, James Foster, Xuechen Li, Harald Oberhauser, and Terry Lyons. Neural SDEs as Infinite-Dimensional GANs. *International Conference on Machine Learning*, 2021.
  - Sang-Yoon Kim and Woochang Lim. Burst synchronization in a scale-free neuronal network with inhibitory spike-timing-dependent plasticity. *Cognitive neurodynamics*, 13(1):53–73, 2019.
  - Xiangwen Kong and Xiangyu Zhang. Understanding masked image modeling via learning occlusion invariant feature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6241–6251, 2023.
  - Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. *International Conference on Artificial Intelligence and Statistics*, 2020.
  - Ying Li, Yanwu Yang, Yuchu Chen, Chenfei Ye, and Ting Ma. Self-supervised learning to unveil brain dysfunctional signatures in brain disorders: Methods and applications. *Health Data Science*, 5:0282, 2025a.
  - Yingzhen Li and Stephan Mandt. Disentangled sequential autoencoder. *arXiv preprint* arXiv:1803.02991, 2018.
  - Yuheng Li, Tianyu Luan, Yizhou Wu, Shaoyan Pan, Yenho Chen, and Xiaofeng Yang. Anatomask: Enhancing medical image segmentation with reconstruction-guided self-masking. In *European Conference on Computer Vision*, pages 146–163. Springer, 2024.
  - Yuheng Li, Yenho Chen, Yuxiang Lai, Jike Zhong, Vanessa Wildman, and Xiaofeng Yang. Medvista3d: Vision-language modeling for reducing diagnostic errors in 3d ct disease detection, understanding and reporting. *arXiv preprint arXiv:2509.03800*, 2025b.
  - Jiayang Liu, Lin Zhong, Jehan Wickramasuriya, and Venu Vasudevan. uwave: Accelerometer-based personalized gesture recognition and its applications. *Pervasive and Mobile Computing*, 5 (6):657–675, 2009.
  - Ziyu Liu, Azadeh Alavi, Minyi Li, and Xiang Zhang. Guidelines for augmentation selection in contrastive learning for time series classification. *arXiv preprint arXiv:2407.09336*, 2024.
  - Francesco Locatello, Ben Poole, Gunnar Rätsch, Bernhard Schölkopf, Olivier Bachem, and Michael Tschannen. Weakly-supervised disentanglement without compromises. In *International conference on machine learning*, pages 6348–6359. PMLR, 2020.

- 594 EN40214341417 Lorenz. E. 1963. deterministic nonperiodic flow. *J. Atmos. Sci*, 20(1):30–41, 1963.
  - Vasilis Z Marmarelis. Nonlinear dynamic modeling of physiological systems. John Wiley & Sons, 2004.
    - TM McKenna, TA McMullen, and MF Shlesinger. The brain as a dynamic physical system. *Neuroscience*, 60(3):587–605, 1994.
    - George Moody. A new method for detecting atrial fibrillation using rr intervals. *Proc. Comput. Cardiol.*, 10:227–230, 1983.
    - Noga Mudrik, Yenho Chen, Eva Yezerets, Christopher J Rozell, and Adam S Charles. Decomposed linear dynamical systems (dlds) for learning the latent components of neural dynamics. *Journal of Machine Learning Research*, 25(59):1–44, 2024.
    - Suraj K Nayak, Arindam Bit, Anilesh Dey, Biswajit Mohapatra, and Kunal Pal. A review on the non-linear dynamical system analysis of electrocardiogram signal. *Journal of healthcare engineering*, 2018(1):6920420, 2018.
    - Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.
    - Chethan Pandarinath, Daniel J O'Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D Stavisky, Jonathan C Kao, Eric M Trautmann, Matthew T Kaufman, Stephen I Ryu, Leigh R Hochberg, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*, 15(10):805–815, 2018.
    - Sam Jean Perochon, Salar Abbaspourazad, Joseph Futoma, Andrew Miller, and Guillermo Sapiro. Time-varying representations of longitudinal biosignals using self-supervised learning. In NeurIPS 2024 Workshop: Self-Supervised Learning-Theory and Practice.
    - JL Reyes-Ortiz, Davide Anguita, Luca Oneto, and Xavier Parra. Smartphone-based recognition of human activities and postural transitions data set. *UCI Machine Learning Repository. School Inf. Comput. Sci. Univ. California at Irvine, Irvine, CA, USA, available online: http://archive.ics. uci. edu/ml/datasets/Smartphone-Based+ Recognition+ of+ Human+ Activities+ and+ Postural+ Transitions, 2015.*
    - Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM international conference on multimodal interaction*, pages 400–408, 2018.
    - Andrew R Sedler and Chethan Pandarinath. Ifads-torch: A modular and extensible implementation of latent factor analysis via dynamical systems. *arXiv preprint arXiv:2309.01230*, 2023.
    - Krishna V Shenoy, Matthew T Kaufman, Maneesh Sahani, and Mark M Churchland. A dynamical systems view of motor preparation: implications for neural prosthetic system design. *Progress in brain research*, 192:33–58, 2011.
    - Arthur Sherman. Dynamical systems theory in physiology. *Journal of General Physiology*, 138(1): 13–19, 2011.
    - Idan Sorin and Michael Tulchinsky. Infinite bifurcations in thomas system. *arXiv preprint arXiv:2408.09525*, 2024.
  - Colin Sparrow. *The Lorenz equations: bifurcations, chaos, and strange attractors*, volume 41. Springer Science & Business Media, 2012.
  - David Sussillo, Rafal Jozefowicz, LF Abbott, and Chethan Pandarinath. Lfads-latent factor analysis via dynamical systems. *arXiv preprint arXiv:1608.06315*, 2016.
    - René Thomas. Deterministic chaos seen in terms of feedback circuits: Analysis, synthesis," labyrinth chaos". *International Journal of Bifurcation and Chaos*, 9(10):1889–1905, 1999.

- Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What makes for good views for contrastive learning? *Advances in neural information processing systems*, 33:6827–6839, 2020.
- Sana Tonekaboni, Danny Eytan, and Anna Goldenberg. Unsupervised representation learning for time series with temporal neighborhood coding. *arXiv* preprint arXiv:2106.00750, 2021.
- Julius Von Kügelgen, Yash Sharma, Luigi Gresele, Wieland Brendel, Bernhard Schölkopf, Michel Besserve, and Francesco Locatello. Self-supervised learning with data augmentations provably isolates content from style. *Advances in neural information processing systems*, 34:16451–16467, 2021.
- Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bousseljot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. Ptb-xl, a large publicly available electrocardiography dataset. *Scientific data*, 7(1):1–15, 2020.
- Jinqiang Wang, Wenxuan Cui, Tao Zhu, Huansheng Ning, and Zhenyu Liu. An improved masking strategy for self-supervised masked reconstruction in human activity recognition. *IEEE Sensors Journal*, 2024.
- Tete Xiao, Xiaolong Wang, Alexei A Efros, and Trevor Darrell. What should not be contrastive in contrastive learning. *arXiv preprint arXiv:2008.05659*, 2020.
- Maxwell A Xu, Alexander Moreno, Hui Wei, Benjamin M Marlin, and James M Rehg. Rebar: Retrieval-based reconstruction for time-series contrastive learning. *arXiv preprint arXiv:2311.00519*, 2023.
- Li Yingzhen and Stephan Mandt. Disentangled sequential autoencoder. In *International Conference on Machine Learning*, pages 5670–5679. PMLR, 2018.
- Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 8980–8987, 2022.
- Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. Self-supervised contrastive pre-training for time series via time-frequency consistency. *Advances in neural information processing systems*, 35:3988–4003, 2022.

# **APPENDIX**

# A PROOF FOR THEOREM 1

**Theorem 1.** Given two time series  $\mathbf{Y}_i$  and  $\mathbf{Y}_j$  independently sampled from the same system (i.e.,  $\mathbf{\Theta}_i = \mathbf{\Theta}_j = \mathbf{\Theta}^{(s)}$ ) under the generative process defined by Eq. 1 and Assumption 1, the minimal set of latent variables shared is the system parameters  $\mathbf{\Theta}^{(s)}$  if and only if all observables from one series is fully masked (i.e.,  $\mathbf{m}_{i,1:W,1:M} = 0$  and  $\mathbf{m}_{j,1:W,1:M} = 1$ ).

*Proof.* Our proof relies on Theorem 1 from Kong and Zhang (2023), which establishes that the minimal set of shared information c in a hierarchical DAG is unique and can be identified using Algorithm 1 presented in their work. We describe the generative model that is considered and summarize Algorithm 1 before applying it in the proof.

**Two-sample Generative Model.** We consider the cross-reconstruction setting with a pair of time series  $(\mathbf{Y}_i, \mathbf{Y}_j)$  for  $i, j \in \mathcal{I}_s$ , generated from the same system such that  $\mathbf{\Theta}_i = \mathbf{\Theta}_j = \mathbf{\Theta}^{(s)}$ . When considering only two samples, the joint distribution in Eq. 1 reduces to the following form,

$$p(\mathbf{Y}_{i}, \mathbf{Y}_{j}, \mathbf{X}_{i}, \mathbf{X}_{j}, \mathbf{\Theta}^{(s)}) = p(\mathbf{\Theta}^{(s)}) \prod_{n=\{i,j\}} p(\mathbf{X}_{n,t_{0}}) \left[ \prod_{k=1}^{W} p(\mathbf{Y}_{n,t_{k}} | \mathbf{X}_{n,t_{k}}) \right] \left[ \prod_{k=2}^{W} p(\mathbf{X}_{n,t_{k}} | \mathbf{X}_{n,t_{k-1}}, \mathbf{\Theta}^{(s)}) \right].$$
(4)

Moreover, we denote masked and unmasked observables respectively as,

$$\mathbf{Y_m} = \bigcup_{k=\{i,j\}} \{\mathbf{Y}_{k,t} \mid \mathbf{m}_{k,t} = 0\} \quad \text{and} \quad \mathbf{Y_{m_c}} = \bigcup_{k=\{i,j\}} \{\mathbf{Y}_{k,t} \mid \mathbf{m}_{k,t} = 1\}. \quad (5)$$

In words,  $Y_m$  and  $Y_{m_c}$  are the sets of masked and unmasked time-points across both samples, respectively.

Algorithm 1 from Kong and Zhang (2023). This approach uses a two-stage procedure for identifying the minimal set of latent variables  $\mathcal C$  shared between masked and unmasked observables in a hierarchical graphical model. In the first stage, called the *selection stage*, all latent variables that are ancestors of both  $Y_m$  and  $Y_{m_c}$  are located. This is done by collecting all parents of the masked observables and adding those that are also ancestors of the unmasked observables to  $\mathcal C$ . The resulting set  $\mathcal C$  contains all latent variables shared between  $Y_m$  and  $Y_{m_c}$ . In the second stage, called the *pruning stage*, each element of  $\mathcal C$  is checked to ensure that no other element in  $\mathcal C$  lies on the directed path between it and  $Y_{m_c}$ . If such a descendant exists, the ancestor node is removed from  $\mathcal C$ . To summarize, given a hierarchical graphical model with masked and unmasked observables  $Y_m$  and  $Y_{m_c}$ , this algorithm returns  $\mathcal C$ , the set of variables with lowest dimension that block all paths between  $Y_m$  and  $Y_{m_c}$ .

Now, we prove each direction of the if-and-only-if condition separately.

If a full sample is masked, then  $\mathcal{C} = \Theta^{(s)}$ . Without loss of generality, we define full-sample mask as  $\mathbf{m}_i = 0$  and  $\mathbf{m}_j = 1$  for  $i, j \in \mathcal{I}_s$ . In this case, any element in  $\mathcal{C}$  must be a common parent of both samples. As illustrated in Fig. 4A, the only parent node that is shared between  $\mathbf{Y}_i$  and  $\mathbf{Y}_j$  is  $\Theta^{(s)}$ , and thus the only set of variables shared between masked and unmasked regions is the system variables. When Algorithm 1 from (Kong and Zhang, 2023) is applied to the graphical model in Eq. 4 under the full-sample masking scheme, it recovers  $\mathcal{C} = \{\Theta^{(s)}\}$ . In the selection stage,  $\Theta^{(s)}$  is the only parent node connecting both samples  $\mathbf{Y}_i$  and  $\mathbf{Y}_j$ . In the pruning stage, nothing is removed since  $\mathcal{C}$  contains only a single element, implying no additional shared latent variables exist as children of  $\Theta^{(s)}$ . Therefore, under full-sample masking, the minimal set of shared latent variables includes only the shared system parameters  $\Theta^{(s)}$ .

If  $c = \Theta^{(s)}$ , then a full sample is masked. We prove the statement via its contrapositive: if a sample is not fully masked (i.e. contains both masked and unmasked observables), then the minimal set of shared latent variables cannot consist solely of the system parameters.

Intuitively, as shown in Fig. 4B, when a sample contains both masked and unmasked time-points, there is always a latent state variable X that serves as the parent node connecting  $\mathbf{Y_m}$  and  $\mathbf{Y_{m_c}}$ . To formalize this, we define a *subsequence mask* as a consecutive region of masked observables, i.e.,  $\mathbf{m}_{i,t_0:t_1}=0$ , where  $t_0,t_1\in 1,\ldots,T$ ,  $t_0\leq t_1$ , and  $t_1-t_0< T$  for a partial mask in a sample of length T. Thus,  $t_0=t_1$  corresponds to masking a single time point at the index  $t_0$ .

There are three possible types of masked subsequence regions: (1) a mask bordering the left edge, or the beginning of the sample, (2) a mask bordering the right edge, or the end of the sample, and (3) a mask in the middle of the sample that is bordered on the left and right by unmasked regions. We apply Algorithm 1 in Kong and Zhang (2023) to each case to determine what minimal set of shared latent variables are recovered. We use  $\mathcal{C}_{\mathrm{subseq}}$  to denote the minimal set of shared latent variables between masked and unmasked regions induced by a subsequence mask.

Case 1 (Left Edge): When the mask borders the left edge of sample  $\mathbf{Y}_i$ , the masked region satisfies  $t_0=1$  and  $t_1 < T$ . During the selection stage, Algorithm 1 retrieves  $\mathcal{C}_{\mathrm{subseq}} = \{\mathbf{\Theta}^{(s)}, \mathbf{X}_{i,t_1}\}$ , since these nodes are ancestors of both masked and unmasked regions. In the pruning stage,  $\mathbf{\Theta}^{(s)}$  is removed because  $\mathbf{X}_{i,t_1}$  lies on the directed path between  $\mathbf{\Theta}^{(s)}$  and the unmasked region. Therefore,  $\mathcal{C}_{\mathrm{subseq}} = \{\mathbf{X}_{i,t_1}\}$ . Intuitively, the latent variable  $\mathbf{X}_{i,t_1}$  serves as the minimal parent connecting the masked and unmasked unmasked regions,  $\mathbf{Y}_{i,1:t_1}$  and  $\mathbf{Y}_{i,t_1:T}$  respectively.

Case 2 (Right Edge): The right edge follows a similar analysis, where the masked region satisfies  $t_0 > 1$  and  $t_1 = T$ . This results in  $\mathcal{C}_{\mathrm{subseq}} = \{ \Theta^{(s)}, \mathbf{X}_{i,t_0} \}$  after the selection stage, and  $\mathcal{C}_{\mathrm{subseq}} = \{ \mathbf{X}_{i,t_0-1} \}$  after the pruning stage. Thus, the latent variable above the  $\mathbf{X}_{i,t_0-1}$  serves as the lowest-level parent connecting unmasked and masked regions,  $\mathbf{Y}_{i,1:t_0}$  and  $\mathbf{Y}_{i,t_0:T}$  respectively.

Case 3 (Middle): When  $t_0 > 1$  and  $t_1 < T$ , the masked region is bordered on both the left and right by unmasked variables. During the selection stage,  $\mathcal{C}_{\mathrm{subseq}} = \{\mathbf{X}_{i,t_0-1}, \mathbf{X}_{i,t_1}, \boldsymbol{\Theta}^{(s)}\}$ , and after pruning,  $\mathcal{C}_{\mathrm{subseq}} = \{\mathbf{X}_{i,t_0-1}, \mathbf{X}_{i,t_1}\}$ , since the latent state variables lie on the directed paths from the system variables to the unmasked regions. Therefore, there are two minimal parent nodes: one at the left boundary,  $\mathbf{X}_{i,t_0-1}$ , above the unmasked variable  $\mathbf{Y}_{i,t_0-1}$ , and one at the right boundary,  $\mathbf{X}_{i,t_1}$ , above the masked variable  $\mathbf{Y}_{i,t_1}$ .

Putting these results together, the minimal shared latent variables induced by a subsequence mask is given by,

$$\mathcal{C}_{\text{subseq}} = \begin{cases} \{\mathbf{X}_{i,t_1}\}, & \text{if } t_0 = 1 \text{ and } t_1 < T, \\ \{\mathbf{X}_{i,t_0-1}\}, & \text{if } t_0 > 1 \text{ and } t_1 = T, \\ \{\mathbf{X}_{i,t_0-1}, \mathbf{X}_{i,t_1}\}, & \text{if } t_0 > 1 \text{ and } t_1 < T. \end{cases}$$
(6)

We can extend this result to arbitrary masks, since any masking configuration over timepoints can be expressed as a union of subsequence masks,  $\mathbf{m}_i = \bigcup_k \mathbf{m}_{i,t_0^{(k)}:t_1^{(k)}}$ . To enforce partial masking, we require that  $\mathbf{m}_{i,t} = 0$  for some t and  $\mathbf{m}_{i,t'} = 1$  for some  $t' \neq t$ . Consequently, for arbitrary masks, the minimal shared latent set is the union over the minimal sets for each subsequence mask  $\mathcal{C} = \bigcup_k \mathcal{C}_{\mathrm{subseq}}^{(k)}$ . Since each  $\mathcal{C}_{\mathrm{subseq}}$  contains only latent state variables, the union  $\mathcal{C}$  does not include the system variables. Therefore, we have shown that under partial masking, the minimal set of shared latent variables does not include the system parameters.

To summarize, applying Algorithm 1 from Kong and Zhang (2023) under a full-sample masking scheme yields  $\mathcal{C} = \{\Theta^{(s)}\}$ , since it is the only variable shared between samples in our graphical model. In contrast, under partial masking, the minimal shared latent set includes only the latent state variables at the boundaries of the masked subsequences, and never the system parameters. This follows directly from our hierarchical structure, as the  $\mathbf{X}$  variables always lie on the directed path from  $\mathbf{\Theta}^{(s)}$  to  $\mathbf{Y}_{\mathbf{m}_c}$ .

Thus, under our two-sample generative model and applying Algorithm 1 from Kong and Zhang (2023), the system parameters appear as the minimal shared latent variables if and only if an entire sample is masked.

# **B** SYNTHETIC DATASET DESCRIPTION

Synthetic time-series trials  $\mathbf{y} \in \mathbb{R}^{T \times M}$  of length T and M measurement dimensions are generated by numerically integrating the Stratonovich SDE,

$$d\mathbf{y}_t = f(\mathbf{y}_t)dt + \tilde{\sigma}d\mathbf{B}_t, \tag{7}$$

with a fixed step size  $\mathrm{d}t=10^{-3}$ , where  $f(\cdot)$  is the dynamics function,  $\mathbf{B}_t$  is multidimensional Brownian noise. To ensure that the noise levels are comparable across different systems, we set the diffusion scale as  $\tilde{\sigma}=\sigma\mathrm{RMS}(\mathbf{Y})$  where  $\sigma$  is a dimensionless noise level and  $\mathrm{RMS}(\mathbf{Y})$  is the root-mean-square amplitude of the noiseless time-series and is estimated empirically through samples  $\mathbf{Y}$  from the system. We consider the following noise levels  $\sigma=\{0,1,3,5\}$  and integrate eq. 7 using torchsde (Li et al., 2020; Kidger et al., 2021).

We generate time series from three dynamical systems: Lorenz, Thomas, and Hindmarsh-Rose. These define strange attractor that produces bounded yet nontrivial dynamics. Importantly, for certain parameter regimes, these systems undergo bifurcations, where changes to parameters induce qualitative changes in the dynamics, thereby altering the statistical properties of the resulting time series. We select systems whose behavior is sensitive to parameter changes, as physiological time series are also generated by nonlinear dynamical systems that are highly sensitive to changes in their underlying parameters. For example, the bursting behavior of a neuron can be triggered or suppressed depending on the inputs it receives (Hindmarsh and Rose, 1984; Kim and Lim, 2019). For each parameter setting and noise level, we generate 20 long time series with  $T=10^5$  time-steps from random initial conditions  $\mathbf{Y}_{i,0} \sim \mathcal{N}(0,I)$  and discard the first 200 steps as a burn-in period to ensure convergence to the attractor manifold. Below, we detail the parameters that we consider for each system.

**Lorenz** (**Lorenz**, **1963**). Although originally derived from atmospheric convection, the Lorenz attractor serves as a conceptual tool for studying physiological dynamics. It exhibits bounded, irregular, and parameter-sensitive behavior, features shared by many biological systems such as heart rate variability (Billman, 2011) and neural activity (Chen et al., 2024; Mudrik et al., 2024; Sussillo et al., 2016). This is the canonical 3D nonlinear attractor used to study chaotic behavior in dynamical systems, with a state-space trajectory that resembles butterfly wings. For this system, M=3, and the dynamics are given by,

$$\frac{d\mathbf{y}}{dt} = \begin{bmatrix} s(y_2 - y_1) \\ y_1(\rho - y_3) - y_2 \\ y_1 y_2 - \beta y_3 \end{bmatrix}$$
 (8)

where  $\mathbf{y} = [y_1, y_2, y_3]^{\top}$ . Following prior work (Sparrow, 2012; Kamiya et al., 2024), we fix  $\beta = 8/3$  and s = 28, and sweep  $\rho$  across the following 10 values:  $\{28, 41, 55, 69, 83, 96, 110, 124, 138, 152\}$ . These values span a range of distinct chaotic regimes.

**Thomas (Thomas, 1999).** The Thomas attractor is a 3D strange attractor that produces cylindrically symmetric time series in state space. For this system, M = 3, and the dynamics are given by

$$\frac{d\mathbf{y}}{dt} = \begin{bmatrix} \sin(y_2) - by_1\\ \sin(y_3) - by_2\\ \sin(y_1) - by_3 \end{bmatrix} \tag{9}$$

where we sweep over  $b \in \{0.025, 0.05, 0.075, 0.1, 0.125, 0.15, 0.175, 0.2, 0.225, 0.25\}$ , corresponding to 10 equally spaced values from 0.025 to 0.25 in increments of 0.025. This range was chosen based on prior work (Sorin and Tulchinsky, 2024), which demonstrates significant changes in system behavior between values of 0 and 0.33.

**Hindmarsh-Rose** (**Hindmarsh and Rose, 1984**). This is a 3D dynamical systems model of neuronal activity that exhibits bursting behavior. For this system, M=3, and the dynamics are given by,

$$\frac{d\mathbf{y}}{dt} = \begin{bmatrix} y_2 - ay_1^3 + by_1^2 - y_3 + I\\ c - dy_1^2 - y_2\\ r[s(y_1 - x_R) - y_3] \end{bmatrix},$$
(10)

and we sweep the external current parameter  $I = \{1, 1.33, 1.66, 2, 2.33, 2.66, 3, 3.33, 3.66, 4\}$ , corresponding to 9 equally spaced values from 1 to 4 in increments of 0.33. This range of parameters is chosen based on prior work (Chen et al., 2023a; Dhamala et al., 2004; Goufo et al., 2020), which shows that the system exhibits different spike–burst behaviors within this region.

Given these generated time-series trials, we construct a dataset for each system, parameter setting, and noise level by combining data from five randomly selected parameter values from each system's grid. By randomly selecting these parameter values, we ensure a range of task difficulty where more challenging datasets involve classifying parameter values that are close together, whereas easier datasets involve classifying parameter values that are farther apart. Importantly, we split each trial into 70:15:15 train, validation, and test splits, and then segment each trial into non-overlapping windows of size W=100. For each system and noise level, we measure the classification accuracy averaged over ten random seeds, accounting for both dataset sampling (classification difficulty) and model initialization. The results reported in Table 1 are the average result for all three systems.

## C BASELINE DESCRIPTION

For contrastive learning and SVAE baselines, we fix the encoder across different training objectives to control for the effects of encoder design and isolate the impact of the pretraining objective. Specifically, we adopt the time series encoder from TS2Vec (Yue et al., 2022), which consists of a 10-layer dilated convolutional network with an embedding size of 320. This setup follows the experimental protocol of (Xu et al., 2023) and report the best available performance from prior work or our own experiments. To obtain a representative embedding for each time window, we apply a global max pooling layer to aggregate features across the temporal dimension. Below, we describe each baseline in more detail.

SimCLR (Chen et al., 2020). SimCLR is a simple augmentation-based method that we adapt for time series data to evaluate the effectiveness of a purely augmentation-driven strategy. Three standard augmentations are applied, each with a 50% probability: scaling, which multiplies the entire time series by a factor drawn from U(0.5, 1.5); shifting, which offsets the time series by a random value in the range [-subsequence size, subsequence size]; and jittering, which adds Gaussian noise with a standard deviation equal to 0.2 times the standard deviation of the dataset.

TS2Vec (Yue et al., 2022). TS2Vec is a competitive time-series contrastive learning method for time series that learns time-stamp-level representations through augmentations. It employs hierarchical contrast, combining instance-level and temporal contrast across multiple resolutions to capture scale-invariant representations within augmented context views. In our experiments, we adopt the dilated convolutional encoder from this method and use the official implementation available at https://github.com/yuezhihan/ts2vec.

**REBAR** (**Xu et al., 2023**). REBAR is a recent time-series contrastive learning method that defines positive pairs using a learned similarity measure. This is accomplished through a cross-attention mechanism that identifies class-specific motifs in one subsequence that can be used to reconstruct another. Subsequences that have the lowest reconstruction error are selected as positive pairs for contrastive learning. We use the implementation provided in the official repository: https://github.com/maxxu05/rebar.

PatchTST (Nie et al., 2022). PatchTST is a transformer model developed for forecasting that uses a patching mechanism, in which consecutive blocks of time points are processed together, and incorporates channel independence, processing each channel separately. It achieves strong performance in forecasting tasks and has been used as a backbone for extracting information about underlying physiological states from biosignals (Geenjaar and Lu, 2025). While the original paper focuses primarily on supervised learning, the model can also be trained in a self-supervised fashion using a masked autoencoding (MAE) objective. In our experiments, we use the Hugging Face implementation from the official repository https://github.com/yuqinie98/PatchTST and train PatchTST in SSL mode, using the CLS token as the summary representation for downstream evaluations.

**TimeMAE** (Cheng et al., 2023). TimeMAE explores the idea of block masking in a MAE framework, adapting it to time-series data. During training, random segments of the input time series are masked, and the unmasked portions are passed through an encoder to produce latent repre-

 sentations. The model also introduces a decoupled autoencoder, where masked and unmasked regions are encoded separately, allowing it to extract transferable information between these regions. A lightweight decoder then reconstructs the masked segments, and the model is trained to minimize the reconstruction error. We use the implementation from the official repository: https://github.com/Mingyue-Cheng/TimeMAE.

LFADS (Sussillo et al., 2016; Sedler and Pandarinath, 2023). Latent Factor Analysis via Dynamical Systems (LFADS) is a deep generative model designed to uncover low-dimensional latent dynamics underlying neural population activity (Pandarinath et al., 2018). During inference, a bidirectional GRU produces an initial condition that serves as a summary representation of a time-series window. This initial condition is then evolved forward by a GRU with a global dynamics function to generate a latent time series, which is linearly projected back into data space to reconstruct the original time series. Although LFADS was originally developed for neural spiking data, its Poisson likelihood can be replaced with a Gaussian likelihood to adapt the framework for continuous-valued time series. In our experiments, we modify the official codebase (https://github.com/arsedler9/lfads-torch) by replacing the BiGRU encoder with the same dilated convolution architecture used in previous baselines, and we find that this modification improves performance on downstream tasks.

**DSVAE** (Li and Mandt, 2018). Disentangled Sequential Variational Autoencoder (DSVAE) (Li and Mandt, 2018) is a generative model originally developed for sequential data (video and audio) and has not previously been applied to physiological signals. We include this baseline in our work since it's generative process resembles PULSE and our results show it is a competitive baseline when applied onto physiological time-series. DSVAE uses a BiLSTM and MLP to infer static and dynamic latent variables, which are then used as initial conditions and inputs to an LSTM for generation. Importantly, DSVAE does not remove irrelevant information, as both system and sample-specific latents are observed and reconstructed jointly. In contrast, PULSE explicitly discards information about noise through the its objective that leverages pseudo-pairs. We use the official implementation<sup>2</sup>, replacing the encoder with dilated convolutions, which improves downstream performance.

**PULSE.** This is our cross-reconstruction based method for physiological time-series SSL proposed in this paper.

**Hyperparameters.** We tune hyperparameters for each method using a random search with a budget of 30 trials over a grid of reasonable values. For trial, we uniformly sample within each of the following grids and select the best setting according to the best validation performance. Training hyperparameters are swept over epochs [50,100,200] and learning rates [0.001,0.0005,0.0001]. Model-specific hyperparameters are varied to capture differences in architecture and training objectives. For PULSE, we sweep over the initial condition encoder hyperparameters, including convolution kernel size [3,5,11], dilation [1,2], and hidden dimensionality [64,128]. For the GRU decoder, we sweep across the number of layers [2,3,4] and hidden dimensionality [64,128]. Finally, for the pseudo-pair construction, we vary the number of samples [1,2,3]. Hyperparameter grids for other baselines are included in the code repository.

### D REAL DATASET DESCRIPTION

For the linear probe experiments, we partition the data into training, validation, and test sets with a 70/15/15 inter-subject splits. Note that the total time of the labeled subsequences may not match the full length of the original recordings, as some portions of the data can be unlabeled. Extracted subsequences are non-overlapping.

HAR (Reyes-Ortiz et al., 2015). The HAR dataset consists of time series data recorded from 30 volunteers of ages 19-48 years with a smartphone (Samsung Galaxy S II) attached at the subjects waist. There are 59 samples of 5-minute long time series that are collected at 50 hz. In our experiment, we use as input 6-channels (3-axis linear accelerometer and 3-axis angular velocity). 6-channels recordings Raw accelerometer and gyroscopic sensor data. Subsequences are 2.56 seconds long (128 time steps) which matches the labels from the original work. collected from smartphones. 6-class classification task with 4,600 subsequences with the following activity class label names and

<sup>2</sup>https://github.com/yatindandi/Disentangled-Sequential-Autoencoder

proportions: walking (17.7 %), walking upstairs (7.6 %), walking downstairs (9.1 %), sitting (18.2 %), standing (20.1 %), and laying (20.1%)

ECG (Moody, 1983). We use data from the MIT-BIH Atrial Fibrillation dataset (Moody, 1983). Since no subsequence length was defined in the original work, we adopt 10-second segments, following both prior analyses of this dataset (Tonekaboni et al., 2021; Xu et al., 2023) and the convention used in ECG classification studies more broadly (Wagner et al., 2020). In total, 76,590 distinct subsequences are extracted from 23 recordings, each lasting approximately 9.25 hours and sampled at 250 Hz with two channels. To further improve computational efficiency, each subsequence is downsampled by a factor of five and produces subsequences with 500 time-steps. Of these, 76,567 subsequences are labeled, with 41.7% corresponding to atrial fibrillation and 58.3% to normal rhythm.

**PPG** (Schmidt et al., 2018). This dataset is constructed from the WESAD dataset. There are 15 recordings each of approximately 87 minutes in duration, corresponding to 334,080 samples collected at 64 Hz from a single channel. From these recordings we extract 1,305 distinct 1-minute subsequences, consistent with the segmentation strategy used in the original work. We improve computational efficiency by downsampling each subsequence by a factor of four, so that each subsequence has 960 time steps. Of these, 666 subsequences are annotated with class labels: baseline (42.7%), stress (24.0%), amusement (12.4%), and meditation (20.9%). All signals are denoised following the procedure described in Heo et al. (2021).

EEG (Kemp et al., 2000). Sleep-EDF (Kemp et al., 2000) contains 39 whole-night electroencephalography (EEG) recordings collected using sleep cassettes from 20 healthy subjects. Following the preprocessing protocol of (Chambon et al., 2018), we use two EEG leads (Fpz-Cz and Pz-Oz) to evaluate pretraining, sampled at 100 Hz and segmented into non-overlapping 30-second intervals (3,000 time steps). This yields a total of 35,424 samples with the following class distribution: Wake (22.9%), N1 (8.89%), N2 (51.30%), N3 (9.92%), and REM (7.00%). The data preprocessing is accomplished with the MNE package (Gramfort et al., 2013). For transfer learning, we follow prior experimental conditions and consider a pretrained model that uses a single EEG lead (Fpz-Cz), segmented into windows of length 200.

For the fine-tuning experiments, we follow the preprocessing and dataset-splitting procedure described in Zhang et al. (2022). The corresponding data are publicly available through the links provided in their repository at https://github.com/mims-harvard/TFC-pretraining.

**Epilepsy.** (Andrzejak et al., 2001) The Epilepsy dataset contains 500 single-channel EEG recordings, each lasting 23.6 seconds. To minimize subject-specific bias, the recordings are divided into 11,500 one-second segments and randomly shuffled, with signals sampled at 178 Hz. The dataset has five labels that capture different conditions or recording locations: eyes open, eyes closed, EEG from healthy regions, EEG from tumor regions, and seizure activity. For our experiments, we reduce the task to binary classification by grouping the first four categories as the negative class and using seizure episodes as the positive class. Fine-tuning is performed on a small labeled subset of 60 samples (30 per class), with 20 additional samples (10 per class) used for validation. The model that achieves the best validation performance is then evaluated on the remaining 11,420 test samples.

Gesture. (Liu et al., 2009) The Gesture dataset captures eight distinct hand movements recorded via accelerometers, with each gesture defined by the path of motion. The gestures include swiping left, right, up, or down; waving in clockwise or counterclockwise circles; tracing a square; and drawing a right arrow. Each gesture is assigned a unique classification label. While the original study reports 4,480 recordings, only 440 samples were available from the UCR Database, yielding a balanced dataset with 55 samples per class. The sampling frequency is not specified in the original paper but is assumed to be 100 Hz. Despite its modest size, the dataset provides enough samples for fine-tuning experiments.

#### E DECLARATION OF LLM USAGE

We used LLMs solely for polishing and revising a base version of the text, for example by converting bullet points into prose, shortening sentences, or improving clarity. LLMs were not used in any other stage of this work, such as information retrieval and discovery, research ideation, or narrative development.