arXiv:2510.09764v3 [cs.LG] 19 Aug 2026

Prototype-based Self-Supervised Multimodal Learning for PPG and Accelerometry Signals

Wanting Mao[1,*], Maxwell A Xu[1,*], Harish Haresamudram[1], Mithun Saha[2], Santosh Kumar[2], James Matthew Rehg[1]
[1]University of Illinois Urbana-Champaign, [2]University of Memphis, [*]Co-first authors
{wmao8,maxu,jrehg}@illinois.edu

Abstract

Modeling multi-modal time-series data is critical for capturing system-level dynamics, particularly in biosignals where modalities such as ECG, PPG, EDA, and accelerometry provide complementary perspectives on interconnected physiological processes. While recent self-supervised learning (SSL) advances have improved unimodal representation learning, existing multi-modal approaches often rely on CLIP-style contrastive objectives that overfit to easily aligned features and misclassify valid cross-modal relationships as negatives, resulting in fragmented and non-generalizable embeddings. To overcome these limitations, we propose ProtoMM, a novel SSL framework that introduces a shared prototype dictionary to anchor heterogeneous modalities in a common embedding space. By clustering representations around shared prototypes rather than explicit negative sampling, our method captures complementary information across modalities and provides a coherent "common language" for physiological signals. In this work, we focus on developing a Pulse-Motion foundation model with ProtoMM and demonstrate that our approach outperforms contrastive-only and prior multimodal SSL methods, achieving the best performance among strong prior baselines while offering additional utility for analyzing and interpreting learned features.

1 Introduction

Digital biomarkers (for stress, physical activity, sleep, etc.) obtained from wearable sensors on smart watches and smartphones, provide unprecedented opportunities to give individuals novel insights into their states of health and wellness throughout their daily life, along with new tools for managing their health-related behaviors [34]. In order to realize this potential, however, it is critical to develop effective models for *multi-modal* time series biosignal data, so that complementary sensing modalities can be leveraged to overcome the ambiguities and noise that are inherent in wearable signals collected in the field environment.

Recently, there has been substantial progress in developing unimodal Foundation Models (FMs) which are pre-trained using large datasets on modalities such as accelerometry [50, 52], ECG [1, 23], and PPG [37, 32]. These models have demonstrated effective generalization to downstream tasks and have established new benchmarks for performance. Building on these successes, recent works have focused on the challenge of how to align multiple signal modalities in pretraining multi-modal FMs, often using CLIP-style contrastive objectives that pull temporally aligned signals together [43, 11, 10, 54]. A key challenge is to ensure that the resulting multimodal embedding captures both *between-modality* information (i.e., features shared across modalities, such as the features of the cardiac cycle that are present in all cardiovascular signals such as ECG, PPG, and ICG) and *within-modality* information (i.e., features that are unique to a single modality, such as the signatures of kinematic motion that are present in accelerometry). When modalities are highly complementary, such as PPG and accelerometry, there is a danger that the alignment process could emphasize between-

modality features at the expense of within-modality features. This is because cardiovascular activity (as captured in the PPG signal) is only indirectly connected to kinematic movement (as captured via accelerometry), e.g. through the increase in cardiac activity which accompanies strenuous exercise. Emphasizing signal alignment could inadvertently discard information which is unique to each modality and critical for downstream tasks such as stress detection.

This paper introduces **ProtoMM**, a novel self-supervised multimodal representation learning strategy, for pre-training foundation models on PPG and accelerometry ("Pulse-Motion") signal modalities.

The goal of ProtoMM is to test the hypothesis that alignment of complementary signal modalities can be facilitated by clustering continuous and noisy biosignals into a shared set of prototypes that capture recurring and semantically meaningful patterns. We hypothesize that these prototypes will be effective in encoding within-modality features and preserving them during the embedding process. ProtoMM achieves this goal by first creating multiple augmented views from each modality. Embeddings from all views, derived using different augmentations of the same modality or from different modalities entirely, are then clustered around prototype vectors. The model is trained using a Multimodal Prototype Prediction loss, where the prototypes probabilities of one view must predict the prototype assignment of another. By enforcing this consistency across all pairs of views, the prototypes become discrete, learnable anchors for the shared latent space for both within-modality and between-modality information.

As a result, ProtoMM provides an explicit prototype-based representation that can be directly inspected and analyzed via prototype assignments. For biosignals, these prototypes offer additional utility beyond representation learning, enabling qualitative and quantitative analysis of physiological patterns. We develop and test ProtoMM for the task of joint multi-modal modeling of PPG and accelerometry data. This is beneficial because PPG can be used to detect the physiological stress response through cardiovascular changes [17], while integrating accelerometer data is essential to disambiguate between responses due to physical activity versus responses caused by psychological stress [39, 40].

We validate the potential of ProtoMM via thorough experimental evaluation that first demonstrates how ProtoMM captures within- and between-modality information. Next, we show how the shared prototype space enables direct analysis of learned semantic patterns. With this, ProtoMM achieves superior performance against strong multimodal self-supervised learning methodologies, and we can qualitatively validate that our prototypes capture morphological similarities with higher-level semantic information. We will release our model weights and a codebase with the full training methodology, architecture, and reproducible evaluation code upon acceptance. The main contributions of this work are:

1. We introduce **ProtoMM**, a novel multimodal self-supervised framework for pulse-motion foundation model, that resolves a key limitation of existing alignment methods. By using a shared set of prototypes and a swapped prediction objective, our model is designed to capture both within-modality (unique) and between-modality (shared) information.
2. We conduct extensive experiments by evaluating its transferability on three downstream datasets across six distinct tasks. We obtain superior performance, showing that ProtoMM consistently outperforms leading multimodal and unimodal baselines.
3. We demonstrate that the explicit nature of our prototype-based learning provides additional utility beyond representation learning, enabling interpretable analysis of physiological and behavioral patterns.

2 Related Work

In recent years, learning useful representations from unlabeled sensor data has become the predominant paradigm, leveraging the ease of wearable sensors to record large quantities of data in naturalistic conditions [5]. Popular approaches include future prediction [29, 15], contrasting between randomly augmented segments [7, 16], probabilistic transformation prediction [36, 52], and reconstruction of randomly masked data [12, 29, 51, 25].

Contrastive Representation Learning for Time-Series Data: Prior work has demonstrated that contrasting randomly transformed windows of sensor data is highly effective, e.g., SimCLR [7].

Similarly, motif-based positive pair generation for contrastive training has shown great promise [49, 50]. Using data from a single modality, these methods essentially learn *within-modality information*.

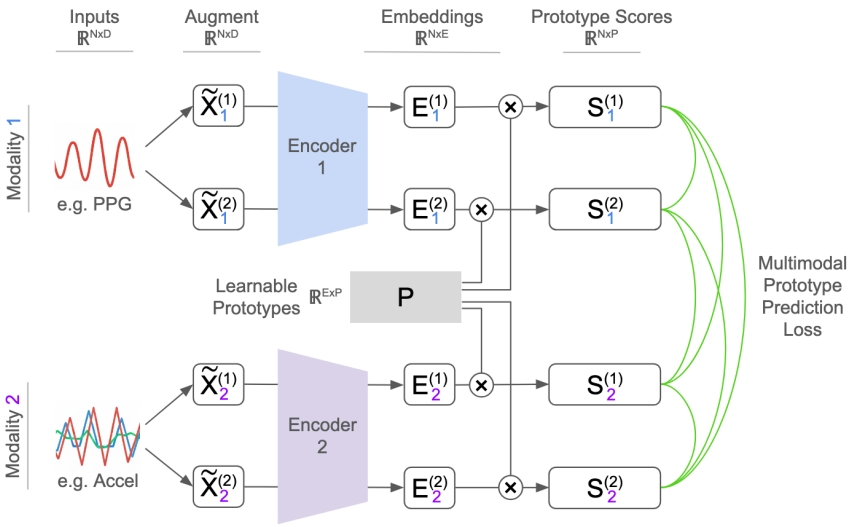

Figure 1: ProtoMM processes augmented segments from multiple modalities (i.e. PPG and Accelerometry) through dedicated encoders to produce the embeddings. The embeddings are then projected onto a shared set of prototype vectors, and the model is trained with a Multimodal Prototype Prediction Loss ($\mathcal{L}_{\text{MPP}}$) that learns to capture both within- and between-modality information without relying on negative sampling.

Alternatively, approaches like ColloSSL [18], CroSSL [10], and COCOA [11] mine positives and negatives *across sensors and modalities.* They evaluate using diverse modalities, including accelerometers, gyroscopes, ECG, EMG, and EDA.

As such, a critical drawback is the non-trivial nature of mining the pairs. More recently, aligning sensor data with natural language descriptions has emerged as an effective option, essentially adopting the CLIP framework [33] for time-series data. Methods such as IMU2CLIP [27], Ts2Act [48], FOCAL [21], and SensorLM [54] have demonstrated the capabilities of such modeling. Further, SLIP [28], which is a combination of the SimCLR and CLIP objectives was evaluated in [13], showcasing improvements. As such, a majority of these methods model only the common *between-modality information.* Our approach—ProtoMM—models both within- and between-modality information. Further, it sidesteps the challenges associated with mining positive/negative pairs by enforcing consistency between cluster assignments of augmented inputs.

Prototype-Based Representation Learning: Instead of instance-based discrimination used in approaches like SimCLR, SwAV [6] jointly clusters the data using prototype vectors, and enforces consistency between soft cluster assignments produced by different augmented inputs. VQ-VAE [45] also employs such vectors, and combines an autoencoder with vector quantization in order to perform online clustering with hard assignment. This setup has been extended to pose data as well [53, 46]. Instead of the autoencoder, VQ-Wav2vec [2, 3] and VQ-CPC [14] use CPC [31] as the base. As such, vector quantization based methods perform hard assignments of the prototype vectors. Our work builds on SwAV, which performs soft cluster assignments, leading to richer expressivity as there can be semantic overlap within the prototype vectors.

3 ProtoMM: Methodology and Design

We introduce ProtoMM, a multimodal self-supervised learning framework for pre-training pulse-motion foundation model, which learns semantically meaningful representations by aligning multimodal time-series data to a shared set of prototypes. Our approach generalizes the swapped assignment prediction mechanism, originally proposed in the unimodal, two-view setting by SwAV [6], to a more general setting involving an arbitrary number of modalities as well as views per modality. By enforcing both within and between-modal consistency, the model learns features that are robust to augmentations while being coherent across different sensor modalities.

3.1 Problem Setup and Notation

We define a multimodal time-series as $\mathbf{X}_t = \{\mathbf{X}_{t,1}, \mathbf{X}_{t,2}, \ldots, \mathbf{X}_{t,M}\}$, where M is the number of sensor modalities and $\mathbf{X}_{t,m} \in \mathbb{R}^{T_m \times C_m}$ represents the data for the m-th modality. The subscript t

indicates that all modal data are temporally aligned to the same window; we omit it for brevity when the context is clear. The dimensions T_m and C_m allow for varying sampling frequencies and channel sizes across modalities.

A data augmentation module, $\mathcal{A}(\cdot)$, is used to generate A distinct views for each modality, creating an augmented set $\{\tilde{\mathbf{X}}_m^{(a)}\}$ for all modalities $m \in \{1, \ldots, M\}$ and views $a \in \{1, \ldots, A\}$. Each augmented view $\tilde{\mathbf{X}}_m^{(a)}$ is then passed through a modality-specific encoder $E_m(\cdot)$ to produce a normalized embedding $\mathbf{E}_m^{(a)} = E_m(\tilde{\mathbf{X}}_m^{(a)})$, where $\mathbf{E}_m^{(a)} \in \mathbb{R}^E$. The resulting set of embeddings can then be aligned using a shared prototype space.

3.2 Shared Prototype Space

To align representations across modalities, we introduce a set of P trainable prototype vectors, organized as columns in a matrix $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_P] \in \mathbb{R}^{E \times P}$. This prototype space is shared across all modalities and training instances, serving as a common set of semantic anchors.

Each embedding $\mathbf{E}_m^{(a)}$ is projected onto the prototypes yielding similarity scores $\mathbf{S} \in \mathbb{R}^P$. To convert these scores into a soft assignment, we use two distinct transformations. First, we compute an assignment target $\mathbf{U}_m^{(a)}$ using the Sinkhorn-Knopp algorithm [8], which enforces an equipartition constraint to prevent mode collapse by ensuring all prototypes are utilized equally across a batch. Second, we compute a probability distribution $\mathbf{V}_m^{(a)}$ using a softmax function with a temperature parameter τ.

$$\mathbf{S}_m^{(a)} = \mathbf{E}_m^{(a)} \mathbf{P} \tag{1}$$

$$\mathbf{U}_m^{(a)} = \mathrm{Sinkhorn}(\mathbf{S}_m^{(a)}) \tag{2}$$

$$\mathbf{V}_m^{(a)} = \mathrm{Softmax}(\mathbf{S}_m^{(a)}/\tau) \tag{3}$$

Now, the assignment target $\mathbf{U}$ and the probability vector $\mathbf{V}$ can be aligned via a cross entropy loss:

$$\ell(\mathbf{U}, \mathbf{V}) = -\mathbf{U} \cdot \log \mathbf{V} \tag{4}$$

3.3 Multimodal Prototype Prediction Loss

The key intuition of our method is that any pair of views originating from the same underlying system, regardless of modality, should be able to predict each other's prototype assignment. This is enforced through a swapped prediction loss comprising two components.

First, the within-modality loss, $\mathcal{L}_{\text{within-mod}}$, enforces consistency between all pairs of augmentations within a modality:

$$\mathcal{L}_{\text{within-mod}} = \sum_{m=1}^{M} \sum_{a=1}^{A} \sum_{b=1, b \neq a}^{A} \ell\Big(\mathbf{U}_m^{(a)}, \mathbf{V}_m^{(b)}\Big) \tag{5}$$

Second, the between-modality loss, $\mathcal{L}_{\text{between-mod}}$, enforces consistency between all pairs from different modalities:

$$\mathcal{L}_{\text{between-mod}} = \sum_{m=1}^{M} \sum_{n=1, n \neq m}^{M} \sum_{a=1}^{A} \sum_{b=1}^{A} \ell\Big(\mathbf{U}_m^{(a)}, \mathbf{V}_n^{(b)}\Big) \tag{6}$$

Our final objective, the Multimodal Prototype Prediction Loss ($\mathcal{L}_{\text{MPP}}$), is a linear combination of these two losses, normalized for stability:

$$\mathcal{L}_{\text{MPP}} = \frac{1}{\binom{A \times M}{2}} (\alpha \mathcal{L}_{\text{within-mod}} + (1 - \alpha) \mathcal{L}_{\text{between-mod}}), \tag{7}$$

The hyperparameter $\alpha \in [0, 1]$ balances the contribution of the two objectives. Setting $\alpha = 1$ reduces the objective to independent, unimodal swapped prediction on each modality found in SwAV [6]. In this work, we set $\alpha = 0.5$ to equally weight both sources of learning. We further show that minimizing $\mathcal{L}_{\text{MPP}}$ is equivalent to maximizing the ELBO of the joint data distribution while the prototypes set act as an information bottleneck to filter out semantically-irrelevant noise. The full derivation is provided in Appendix A.11.

4 Experimental Design

In Section 4.1, we first detail our large-scale pre-training dataset, along with the architecture and training settings. Then, in Section 4.2, we describe the unimodal and multimodal baselines. Finally, in Section 4.3, we discuss our downstream datasets and tasks, as well as the evaluation procedure.

4.1 Pre-training Setup

Our pre-training data comprises the initial 10 days of data from the large-scale Mobile Open Observation of Daily Stressors (MOODS) study [30]. It contains 794,872 synchronized 30-second segments of accelerometer and PPG signals from 122 participants (39 men, 77 women, 6 non-binary; mean age 38±13 years). The data was collected "in-the-wild" from a wrist-worn WearOS smartwatch as participants went about their daily lives without specific instructions. Further details on the study design are available in [30].

To ensure fair comparison, all models and modalities use the same encoder architecture and a consistent set of augmentations specifically for time-series [42]. The only architectural difference is the number of input channels for the initial convolution layer: 3 for accelerometer data, 1 for PPG, and 4 for the unimodal methods that concatenate inputs. In Table 1, the "Input" column denotes early-fusion models with concatenated inputs as P+A and multimodal models with separate encoders as P|A.

We utilize the Adam optimizer [19] with a learning rate of 10^{-5}, no weight decay, and batch size of 256. Models are trained for 100 epochs or 96 hours (whichever comes first) on an NVIDIA L40S GPU. Both training and validation loss are calculated per batch but aggregated at epoch level, and the checkpoint with lowest validation loss is used for all downstream experiments. More details are in Appendix A.4.

4.2 Benchmarks

We compare ProtoMM against a comprehensive suite of self-supervised baselines, all trained from scratch using the identical pre-training setup. Shared or modality-specific projection heads follow each baseline's original specifications. We categorize these baselines into three primary groups.

General Multimodal Alignment methods: General frameworks for cross-modal representation learning. This includes *CLIP* [33, 43], which employs a pairwise contrastive loss to align temporally-paired signals, and *SLIP* [28], which extends this approach by combining SimCLR and CLIP objectives to capture both within- and between-modality interactions.

Time-Series Multimodal Alignment methods: Explicitly designed for multi-sensor data. This consists of *COCOA* [11], which contrasts temporally aligned multi-sensor positives against misaligned within-sensor negatives; *CroSSL* [10], which leverages latent masking and cross-modal aggregation; and *FOCAL* [21], which factorizes representations into orthogonal between- and within-modal subspaces using temporal and spectral augmentations.

Unimodal Self-supervised Learning methods: These methods are adapted for multimodal input through simple concatenation. We evaluate standard *SimCLR* [7], a unimodal contrastive learning framework utilizing augmented versions of input data, and our own unimodal ablation, *ProtoMM Within-Mod* (equivalent to SwAV [6]), which only performs within-modal learning. We also include advanced time-series specific approaches such as *RelCon* [50], which uses a learnable motif-based distance metric to assign soft positive pairs, and the momentum-based framework by *Abbaspourazad et al.* [1], which mines positive pairs from segments of the same subject.

4.3 Downstream Evaluation Setup

For evaluation, we use three datasets that contain synchronous PPG and accelerometer data, along with annotations: MOODS (stress detection, activity recognition) [30], WESAD (stress detection) [38], and PPG-DaLiA (activity recognition, instantaneous heart rate prediction) [35]. Full dataset details are provided in Appendix A.5. All signals are resampled to 50 Hz to match the sampling frequency in pre-training dataset. Following previous works [50, 1, 29, 37], after pre-training is complete, we freeze the model encoders and train a linear probe on the downstream tasks. For multimodal models,

Table 1: **Linear-probe evaluation of unimodal and multimodal baselines.** ProtoMM achieves best overall performance which indicates that grounding the alignment in a shared, discrete prototype space is a more effective mechanism for learning both shared and unique features simultaneously.

			MOODS				WESAD				PPG-DaLiA			
			Stress (2)		Activity (2)		Stress (4)		Stress (2)		Activity (9)		HR (R)	
Model	In	Out	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↓MAE	↑ R^2
SimCLR	A	A	0.477±.006	0.598±.006	0.861±.004	0.925±.002	0.557±.053	0.629±.052	0.793±.055	0.836±.044	0.559±.005	0.560±.004	15.37±.10	0.101±.015
SimCLR	P	P	0.527±.006	0.607±.006	0.655±.005	0.857±.002	0.348±.034	0.517±.053	0.692±.056	0.699±.054	0.302±.004	0.399±.004	**8.14±.07**	**0.670±.008**
PMM WMod	A	A	0.464±.006	0.604±.006	0.857±.004	0.924±.002	0.533±.057	0.584±.053	0.675±.060	0.726±.053	0.586±.005	0.577±.004	15.45±.10	0.097±.015
PMM WMod	P	P	0.469±.005	0.611±.005	0.595±.005	0.848±.003	0.522±.055	0.607±.053	0.847±.050	0.877±.039	0.285±.004	0.397±.004	8.29±.07	**0.670±.007**
SimCLR	P+A	P+A	0.488±.006	0.601±.006	0.849±.004	0.919±.002	0.445±.048	0.596±.053	0.721±.056	0.753±.051	0.542±.005	0.536±.004	9.91±.08	0.534±.009
PMM WMod	P+A	P+A	0.467±.006	0.600±.006	0.836±.004	0.913±.002	0.486±.054	0.562±.055	0.753±.056	0.795±.047	0.543±.005	0.547±.005	15.07±.09	0.155±.014
Abbaspourazad et al.	P+A	P+A	0.445±.005	0.614±.005	0.681±.005	0.865±.002	0.417±.044	0.551±.056	0.668±.062	0.753±.050	0.444±.005	0.453±.004	15.38±.10	0.103±.014
RelCon	P+A	P+A	0.520±.006	0.589±.005	0.748±.004	0.880±.002	0.573±.052	0.640±.051	**0.910±.035**	**0.918±.031**	0.433±.005	0.461±.005	9.63±.08	0.578±.009
CLIP	P|A	P+A	0.524±.006	0.578±.006	0.869±.003	0.929±.002	0.496±.048	0.618±.053	**0.910±.034**	**0.918±.031**	0.640±.005	0.624±.004	9.37±.07	0.609±.008
COCOA	P|A	P+A	0.520±.006	0.579±.005	0.858±.003	0.924±.002	0.508±.050	0.607±.053	0.778±.053	0.808±.046	0.620±.005	0.602±.004	9.43±.07	0.615±.008
CroSSL	P|A	P+A	0.461±.005	**0.622±.006**	0.751±.004	0.882±.002	0.430±.045	0.494±.052	0.783±.053	0.808±.047	0.553±.005	0.540±.004	11.47±.08	0.464±.010
FOCAL	P|A	P+A	0.510±.006	0.588±.006	0.853±.003	0.921±.002	0.596±.049	0.708±.049	0.888±.041	0.904±.035	0.607±.005	0.606±.004	11.66±.08	0.449±.010
SLIP	P|A	P+A	0.524±.006	0.586±.006	0.867±.003	0.929±.002	0.500±.039	0.652±.051	0.848±.046	0.863±.040	0.633±.004	0.625±.004	8.89±.08	0.634±.008
ProtoMM	P|A	P+A	**0.532±.006**	0.591±.006	**0.872±.003**	**0.932±.002**	**0.622±.050**	**0.719±.048**	**0.910±.035**	**0.918±.031**	**0.656±.004**	**0.638±.004**	8.74±.07	0.648±.008

KEY: Encoder (In)put, Final Embedding (Out)put; (A)ccel, (P)PG; + designates concatenation, | designates separate encoders

we generally will concatenate the embeddings from the PPG and accelerometer encoders to form the final representation. This is indicated by P+A in the "Out" column in our Results Table 1 and 3. For unimodal baselines, we use the single embedding directly. This is indicated by a P in the Out column in our Results Table 1. We report macro F1-score and accuracy for classification tasks (stress, activity) and mean absolute error (MAE) and R^2 for the regression task (heart rate prediction).

5 Discussion

In this section, we organize the empirical analysis into three parts. First, we evaluate ProtoMM in the main pulse-motion setting using PPG and accelerometry signals, where we compare against a broad set of multimodal benchmarks. Then, we analyze the prototypes' utility and visualize the learned prototype space to show how they capture semantic meaning and specific morphologies. Finally, we study how ProtoMM learns specific within- and between-modality information, as well as how it generalizes to new multimodal time-series domains.

5.1 Analysis of Benchmark Results

ProtoMM Achieves the Strongest Performance. Table 1 presents comparisons against all baseline methods, which show that ProtoMM achieves the strongest results. It achieves the best overall performance in 4/6 downstream tasks and outperforms all other multimodal methods in 5/6 tasks.

Within-modality and Between-modality Information. Out of the multimodal baselines, CLIP, COCOA, and CroSSL focusing on modeling only the between-modal information, whereas SLIP, FOCAL and our ProtoMM method are the ones that explicitly model both between and within-modal information. However, interestingly, the models that model both do not uniformly outperform the models that model only between-modal information. Despite SLIP augmenting the CLIP objective with an additional within-modal loss, their overall performance is comparable as the 2nd best models after ProtoMM. This suggests that standard contrastive losses may struggle to balance the two objectives, potentially over-emphasizing the more difficult between-modal alignment. ProtoMM's success indicates that grounding the alignment in a shared, discrete prototype space is a more effective mechanism for learning both shared and unique features simultaneously.

Multimodality vs Unimodality. The results show a clear trend where multimodal models outperform their unimodal counterparts, particularly on complex classification tasks. This highlights the value of alignment: one modality provides essential context that the other lacks. For example, HR information from PPG can help disambiguate activities with similar motion profiles from accelerometry, such as ascending versus descending stairs classes in PPG DaLiA activity classification or how research has shown that accelerometry and PPG signals can be used together for stress prediction [39, 47]. The one exception to this trend is in the HR regression tasks. This highlights a key nuance in multimodal learning. HR is a metric that can be derived directly from the raw PPG waveform, and accelerometry signal gives little to no information on it. Therefore, the inclusion of accelerometry within the embedding model only serves to further obfuscate the final embedding, such that multimodal models generally do worse. However, ProtoMM is able to more effectively disentangle the modalities, to better preserve within-modal PPG-specific information, achieving the best multimodal performance.

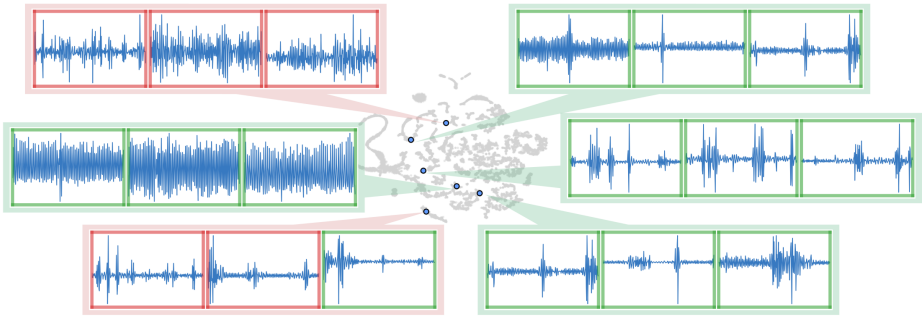

(a) Top-three nearest **PPG** time-series for representative prototype centroids.

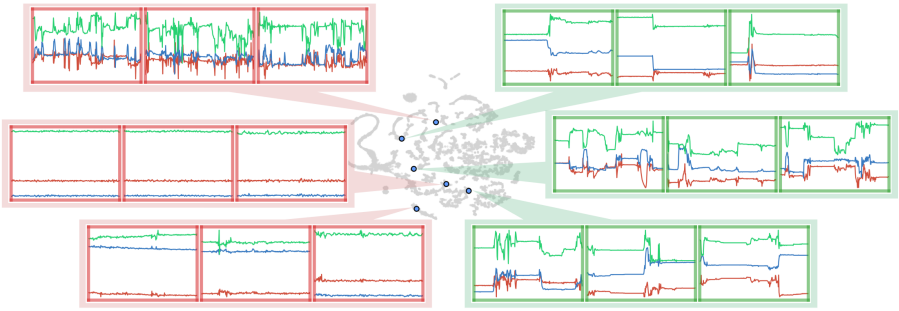

(b) Top-three nearest **accelerometer** time-series for representative prototype centroids.

Figure 2: t-SNE of learned prototypes (gray), with k-means centroids (blue) and their top-three nearest time-series from two modalities: **PPG** (top) and **accelerometer** (bottom). Panel borders denote ground-truth labels (green = Unstressed, red = Stressed). Neighbors around each prototype show consistent labels and morphology; the prototypes capture distinct patterns. The shared prototypes semantically anchor the representation, unifying observable patterns and underlying physiological context to expose latent states.

Multi-modal SSL Models are Important. Finally, the table shows the unimodal models that simply concatenate the raw signals perform poorly. Within Table 1, they are marked as P+A in the In column. This finding demonstrates that despite both being biosignals, a naive concatenation of PPG and accelerometry is insufficient, and modality-specific encoders are essential for learning how to capture PPG-specific and Accelerometry-specific features from the raw sensor data.

5.2 Understanding the Prototypes in <u>Proto</u>MM

Prototypes Explicitly Improve Performance. SLIP serves as the prototype-free analogue to ProtoMM. We utilize identical architecture and augmentation set for both models, making them equivalent up to the final embedding layer that produces $\mathbf{E}_m^{(a)}$. Then, ProtoMM projects these embeddings onto a learned set of prototypes before applying a swapped prediction loss across between- and within-modality pairs, but SLIP directly applies a contrastive (NT-Xent) loss to the embeddings themselves across all between- and within-modality pairs. Consequently, this comparison enables us to quantify the contribution of the prototype mechanism.

The results in Table 1 confirm the advantages of our prototype-based method. It demonstrates superior performance over its direct prototype-free analogue, SLIP, on every metric across all six downstream tasks. The performance gains are particularly pronounced on the WESAD dataset, where ProtoMM improves the F1-score for 4-class stress detection by a significant 24.6% (0.622 vs. 0.500) and for binary stress detection by over 7% (0.910 vs. 0.848).

Interestingly, performance gains remain isolated to the multimodal setting. In unimodal settings, SimCLR consistently outperforms its prototype-based counterpart, ProtoMM W-Mod. This aligns with prior work in unimodal time-series self-supervision [24]. We hypothesize this is for two reasons: first, prototypes may act as a shared, discretized vocabulary that provides a common language

Table 2: **Integrating multimodal information into a unimodal embedding improves performance.** ProtoMM rows show the performance of an unimodal embedding that was pre-trained multimodally (with both PPG and Accelerometer), while ProtoMM Within-Modality rows show the performance of an embedding pre-trained in isolation on just that single sensor.

				MOODS				WESAD				PPG-DaLiA			
				Stress (2)		Activity (2)		Stress (4)		Stress (2)		Activity (9)		HR (R)	
α	Model	In	Out	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↓MAE	↑ R^2
.5	**ProtoMM**	P|A	A	$\mathbf{0.496}_{\pm.006}$	$0.597_{\pm.006}$	$\mathbf{0.872}_{\pm.003}$	$\mathbf{0.931}_{\pm.002}$	$\mathbf{0.620}_{\pm.060}$	$\mathbf{0.663}_{\pm.051}$	$\mathbf{0.837}_{\pm.048}$	$\mathbf{0.863}_{\pm.041}$	$\mathbf{0.615}_{\pm.004}$	$\mathbf{0.597}_{\pm.004}$	$\mathbf{15.0}_{\pm.10}$	$\mathbf{0.152}_{\pm.014}$
1	PMM WMod	A	A	$0.464_{\pm.006}$	$\mathbf{0.604}_{\pm.006}$	$0.857_{\pm.004}$	$0.924_{\pm.002}$	$0.533_{\pm.057}$	$0.584_{\pm.053}$	$0.675_{\pm.060}$	$0.726_{\pm.053}$	$0.586_{\pm.005}$	$0.577_{\pm.004}$	$15.5_{\pm.10}$	$0.097_{\pm.015}$
.5	**ProtoMM**	P|A	P	$\mathbf{0.541}_{\pm.006}$	$\mathbf{0.613}_{\pm.006}$	$\mathbf{0.747}_{\pm.004}$	$\mathbf{0.882}_{\pm.002}$	$0.518_{\pm.038}$	$\mathbf{0.607}_{\pm.052}$	$\mathbf{0.855}_{\pm.043}$	$0.863_{\pm.040}$	$\mathbf{0.332}_{\pm.005}$	$\mathbf{0.426}_{\pm.005}$	$\mathbf{8.16}_{\pm.07}$	$\mathbf{0.686}_{\pm.007}$
1	PMM WMod	P	P	$0.469_{\pm.005}$	$0.611_{\pm.005}$	$0.595_{\pm.005}$	$0.848_{\pm.003}$	$\mathbf{0.522}_{\pm.055}$	$\mathbf{0.607}_{\pm.053}$	$0.847_{\pm.050}$	$\mathbf{0.877}_{\pm.039}$	$0.285_{\pm.004}$	$0.397_{\pm.004}$	$8.29_{\pm.07}$	$0.670_{\pm.007}$

KEY: Encoder (In)put, Final Embedding (Out)put; (A)ccel, (P)PG; + designates concatenation, | designates separate encoders

to translate between disparate data streams like PPG and accelerometry. In a unimodal setting, the augmented views already reside on the same manifold, rendering this translation mechanism redundant. Second, the discrete bottleneck acts as a regularizer designed to filter out non-shared, modality-specific noise. In the unimodal case, where such independent noise sources are absent, this discretization acts effectively as a low-pass filter, discarding fine-grained continuous features that SimCLR preserves for richer instance discrimination.

Visualizing Prototypes. To demonstrate that ProtoMM learns an explicit prototype-based representation that captures recurring and semantically meaningful signal patterns, we conduct a qualitative analysis on the WESAD dataset. We cluster the learned prototypes with k-means clustering (k=15), then for a given learned prototype centroid, we identify the accelerometer and PPG segments with highest cosine similarity in the representation space. Figure 2a and 2b show a t-SNE projection of all prototypes, with representative examples highlighted alongside their top-3 nearest neighbors.

These neighbors exhibit both label consistency (stressed vs. unstressed) and coherent temporal dynamics that correspond to distinct physiological patterns (e.g., stable low-amplitude patterns, sharp oscillations). This suggests that the shared prototype vectors are structured in the representation space in clusters that exhibit both label consistency and coherent physiological patterns. Rather than operating as a black box, ProtoMM's prototypes function as semantically meaningful anchors that capture both morphological signal characteristics and higher-level physiological contexts, offering a tangible insights into the latent physiological state. Appendix A.7 further shows that the learned prototype space contains both modality-specific and modality-shared prototypes, providing additional evidence that ProtoMM preserves complementary sensor information while aligning shared structure.

5.3 Understanding the Multi-Modality in Proto<u>MM</u>

Between-Modal Knowledge Transfer to the Other Modality. We would like to further explore how effective ProtoMM integrates Between-modal information by training ProtoMM normally, to capture both between-modal and within-modal information, and then instead of concatenating the outputs of each modality's embedding for evaluation, just using one modality's embedding at a time, show by the 1st and 3rd rows in Table 2. This will help us investigate whether the modality's trained embedding captures more information than if it was trained independently. Therefore, the baseline is training each modality independently with a unimodal ProtoMM Within-Mod ($\alpha = 1$) with only one modality as input, shown by the 2nd and 4th rows in Table 2.

Balancing Within- and Between-Modality Objectives. The core advantage of ProtoMM lies in its ability to simultaneously and effectively capture between- and within-modality information through a unified objective, $\mathcal{L}_{\text{MPP}}$. We validate this design choice by modifying the loss weighting parameter to be $\alpha = 1$, such that only within-modality ($\mathcal{L}_{\text{within-mod}}$) information is learned or to be $\alpha = 0$, such that only between-modality ($\mathcal{L}_{\text{between-mod}}$) information is learned.

Table 3 shows that performance is generally higher when α is not 0 or 1 (best at $\alpha = 0.5$ and 0.67) , demonstrating that successful multimodal integration requires explicitly modeling both the unique contributions of each modality and their synergistic interactions. The balanced objective enables ProtoMM to learn representations that are simultaneously invariant to modality-specific augmentations while being semantically aligned across modalities.

Table 3: **Balancing within- and between-modal objectives is necessary.** ProtoMM achieves the best performance at α=0.5 and 0.67, showing that the unified MPP objective improves performance beyond a strict between-modality (α=0) or within-modality (α=1) objective.

			α					
Dataset	Task	Metric	0 (BMod)	0.25	0.5 (ProtoMM)	0.67	0.75	1 (WMod)
MOODS	Stress (2)	↑F1	$0.461_{\pm.005}$	$0.463_{\pm.005}$	$0.532_{\pm.006}$	$0.522_{\pm.006}$	$\mathbf{0.537_{\pm.006}}$	$0.497_{\pm.006}$
		↑Acc	$0.604_{\pm.006}$	$\mathbf{0.608_{\pm.005}}$	$0.591_{\pm.006}$	$0.586_{\pm.006}$	$0.600_{\pm.006}$	$0.591_{\pm.006}$
	Activity (2)	↑F1	$0.719_{\pm.005}$	$0.724_{\pm.005}$	$\mathbf{0.872_{\pm.003}}$	$0.869_{\pm.003}$	$0.871_{\pm.003}$	$0.855_{\pm.004}$
		↑Acc	$0.874_{\pm.002}$	$0.876_{\pm.002}$	$\mathbf{0.932_{\pm.002}}$	$0.929_{\pm.002}$	$0.931_{\pm.002}$	$0.923_{\pm.002}$
WESAD	Stress (4)	↑F1	$0.578_{\pm.057}$	$0.589_{\pm.058}$	$\mathbf{0.622_{\pm.050}}$	$0.577_{\pm.054}$	$0.526_{\pm.047}$	$0.612_{\pm.055}$
		↑Acc	$0.663_{\pm.050}$	$0.685_{\pm.052}$	$\mathbf{0.719_{\pm.048}}$	$0.674_{\pm.050}$	$0.640_{\pm.051}$	$0.685_{\pm.050}$
	Stress (2)	↑F1	$0.783_{\pm.052}$	$0.786_{\pm.054}$	$0.910_{\pm.035}$	$\mathbf{0.938_{\pm.031}}$	$0.800_{\pm.051}$	$0.858_{\pm.049}$
		↑Acc	$0.808_{\pm.045}$	$0.822_{\pm.044}$	$0.918_{\pm.031}$	$\mathbf{0.945_{\pm.027}}$	$0.822_{\pm.045}$	$0.890_{\pm.037}$
PPG-DaLiA	Activity (9)	↑F1	$0.496_{\pm.005}$	$0.485_{\pm.005}$	$0.656_{\pm.004}$	$0.657_{\pm.004}$	$\mathbf{0.662_{\pm.004}}$	$0.618_{\pm.005}$
		↑Acc	$0.502_{\pm.004}$	$0.493_{\pm.004}$	$0.638_{\pm.004}$	$\mathbf{0.651_{\pm.004}}$	$0.648_{\pm.004}$	$0.603_{\pm.004}$
	HR (R)	↓MAE	$10.45_{\pm.08}$	$10.76_{\pm.08}$	$8.74_{\pm.07}$	$\mathbf{8.67_{\pm.07}}$	$9.08_{\pm.08}$	$9.14_{\pm.08}$
		↑ R^2	$0.536_{\pm.009}$	$0.515_{\pm.009}$	$0.648_{\pm.008}$	$\mathbf{0.654_{\pm.007}}$	$0.618_{\pm.008}$	$0.611_{\pm.008}$

Table 4: **Generalization to new, additional modalities.** Models are trained and evaluated on MOD and RealWorld-HAR. MOD contains the two modalities of audio and seismic signals, and RealWorld-HAR contains four modalities of accelerometry, gyroscope, magnetometer, and light signals. ProtoMM achieves the best performance under both linear and MLP probe eval, indicating generalizability to new modalities beyond the original PPG + accelerometry setting and to 2+ modalities.

	MOD				RealWorld-HAR			
	Linear		MLP		Linear		MLP	
Model	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc
CLIP	$0.653\pm.018$	$0.668\pm.018$	$0.639\pm.018$	$0.659\pm.018$	$0.714\pm.003$	$0.696\pm.003$	$0.686\pm.004$	$0.682\pm.004$
COCOA	$0.652\pm.018$	$0.668\pm.017$	$0.625\pm.017$	$0.647\pm.017$	$0.701\pm.004$	$0.694\pm.003$	$0.681\pm.004$	$0.682\pm.003$
CroSSL	$0.635\pm.018$	$0.651\pm.018$	$0.637\pm.018$	$0.655\pm.018$	$0.611\pm.004$	$0.619\pm.004$	$0.584\pm.004$	$0.608\pm.004$
FOCAL	$0.591\pm.018$	$0.596\pm.018$	$0.538\pm.018$	$0.566\pm.018$	$0.659\pm.004$	$0.651\pm.003$	$0.645\pm.004$	$0.643\pm.003$
SLIP	$0.710\pm.018$	$0.725\pm.017$	$0.659\pm.018$	$0.690\pm.017$	$0.704\pm.003$	$0.690\pm.003$	$0.629\pm.004$	$0.633\pm.004$
ProtoMM	$\mathbf{0.740\pm.016}$	$\mathbf{0.749\pm.016}$	$\mathbf{0.722\pm.016}$	$\mathbf{0.733\pm.016}$	$\mathbf{0.732\pm.003}$	$\mathbf{0.712\pm.003}$	$\mathbf{0.717\pm.003}$	$\mathbf{0.701\pm.003}$

As shown in Table 2, both encoders (i.e., for the accelerometer and PPG) outperform their unimodal ProtoMM Within-Mod counterparts on nearly all tasks, despite being evaluated without the full multimodal embedding. This indicates that the shared prototype space encourages each encoder to develop representations that are semantically aligned with the broader physiological context across modalities, not just its own signal characteristics, which results in more robust and informative features even when deployed as a unimodal embedding.

Generalization to New Modalities. To further demonstrate the broad applicability of our method across multimodal time-series domains, we conduct an additional set of experiments on two multimodal time-series benchmarks. Moving Object Detection (MOD) [21] contains acoustic and seismic signals for vehicle classification, while RealWorld-HAR [41] contains four synchronized wearable modalities: accelerometer, gyroscope, magnetometer, and light sensors. As shown in Table 4, ProtoMM achieves the best performance on both datasets under both linear and MLP evaluation. These results suggest ProtoMM's ability to transfer to non-physiological sensing modalities and scale beyond two-modality pairs. Full experimental details are provided in Appendix A.6.

6 Conclusion

In this paper, we presented *ProtoMM*, a prototype-based multimodal framework for self-supervised learning on time-series data to pre-train a pulse-motion foundation model. By leveraging a shared prototype space, it aligns embeddings from different modalities and uncovers common latent physiological states. The learned prototypes act as cluster centers, structuring the representation space into coherent groups that reflect both morphological similarities and higher-level semantic consistency. Importantly, this enhances interpretability, as individual prototypes correspond to meaningful physiological patterns. Comprehensive experiments across three datasets and six downstream tasks demonstrate that ProtoMM consistently outperforms a broad set of strong baselines. Beyond contrastive or masking-based approaches, ProtoMM introduces a prototype-based swapped prediction objective, offering a new perspective on representation learning for multimodal time-series.

References

[1] Salar Abbaspourazad, Oussama Elachqar, Andrew C Miller, Saba Emrani, Udhyakumar Nallasamy, and Ian Shapiro. Large-scale training of foundation models for wearable biosignals. *arXiv preprint arXiv:2312.05409*, 2023.

[2] Alexei Baevski, Steffen Schneider, and Michael Auli. vq-wav2vec: Self-supervised learning of discrete speech representations. *arXiv preprint arXiv:1910.05453*, 2019.

[3] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460, 2020.

[4] biosignalsplux. Biosignalsplux Explorer User Manual. `https://bio-medical.com/media/support/biosignalsplux_explorer_user_manual_v.1.0.pdf`, 2019.

[5] Clare Bycroft, Colin Freeman, Desislava Petkova, Gavin Band, Lloyd T Elliott, Kevin Sharp, Allan Motyer, Damjan Vukcevic, Olivier Delaneau, Jared O'Connell, et al. The uk biobank resource with deep phenotyping and genomic data. *Nature*, 562(7726):203–209, 2018.

[6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020.

[7] 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.

[8] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

[9] Kamana Dahal, Brian Bogue-Jimenez, and Ana Doblas. Global stress detection framework combining a reduced set of hrv features and random forest model. *Sensors*, 23(11):5220, 2023.

[10] Shohreh Deldari, Dimitris Spathis, Mohammad Malekzadeh, Fahim Kawsar, Flora D Salim, and Akhil Mathur. Crossl: Cross-modal self-supervised learning for time-series through latent masking. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 152–160, 2024.

[11] Shohreh Deldari, Hao Xue, Aaqib Saeed, Daniel V Smith, and Flora D Salim. Cocoa: Cross modality contrastive learning for sensor data. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(3):1–28, 2022.

[12] Harish Haresamudram, Apoorva Beedu, Varun Agrawal, Patrick L Grady, Irfan Essa, Judy Hoffman, and Thomas Plötz. Masked reconstruction based self-supervision for human activity recognition. In *Proceedings of the 2020 ACM International Symposium on Wearable Computers*, pages 45–49, 2020.

[13] Harish Haresamudram, Apoorva Beedu, Mashfiqui Rabbi, Sankalita Saha, Irfan Essa, and Thomas Ploetz. Limitations in employing natural language supervision for sensor-based human activity recognition-and ways to overcome them. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 273–281, 2025.

[14] Harish Haresamudram, Irfan Essa, and Thomas Ploetz. Towards learning discrete representations via self-supervision for wearables-based human activity recognition. *Sensors*, 24(4):1238, 2024.

[15] Harish Haresamudram, Irfan Essa, and Thomas Plötz. Contrastive predictive coding for human activity recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(2):1–26, 2021.

[16] Harish Haresamudram, Irfan Essa, and Thomas Plötz. Assessing the state of self-supervised human activity recognition using wearables. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 6(3), sep 2022.

[17] Anice Jahanjoo, Nima TaheriNejad, and Amin Aminifar. High-accuracy stress detection using wrist-worn ppg sensors. In *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5, 2024.

[18] Yash Jain, Chi Ian Tang, Chulhong Min, Fahim Kawsar, and Akhil Mathur. Collossl: Collaborative self-supervised learning for human activity recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(1):1–28, 2022.

[19] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[20] Lucas Lange, Nils Wenzlitschke, and Erhard Rahm. Generating synthetic health sensor data for privacy-preserving wearable stress detection. *Sensors*, 24(10):3052, 2024.

[21] Shengzhong Liu, Tomoyoshi Kimura, Dongxin Liu, Ruijie Wang, Jinyang Li, Suhas Diggavi, Mani Srivastava, and Tarek Abdelzaher. Focal: Contrastive learning for multimodal time-series sensing signals in factorized orthogonal latent space. *Advances in Neural Information Processing Systems*, 36:47309–47338, 2023.

[22] Cameron McCarthy, Nikhilesh Pradhan, Calum Redpath, and Andy Adler. Validation of the empatica e4 wristband. In *2016 IEEE EMBS international student conference (ISC)*, pages 1–4. IEEE, 2016.

[23] Kaden McKeen, Laura Oliva, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. Ecg-fm: An open electrocardiogram foundation model, 2024.

[24] Qianwen Meng, Hangwei Qian, Yong Liu, Lizhen Cui, Yonghui Xu, and Zhiqi Shen. Mhccl: masked hierarchical cluster-wise contrastive learning for multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 9153–9161, 2023.

[25] Shenghuan Miao, Ling Chen, and Rong Hu. Spatial-temporal masked autoencoder for multi-device wearable human activity recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 7(4):1–25, 2024.

[26] Varun Mishra, Tian Hao, Si Sun, Kimberly N Walter, Marion J Ball, Ching-Hua Chen, and Xinxin Zhu. Investigating the role of context in perceived stress detection in the wild. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, pages 1708–1716, 2018.

[27] Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Alireza Dirafzoon, Aparajita Saraf, Amy Bearman, and Babak Damavandi. Imu2clip: Multimodal contrastive learning for imu motion sensors from egocentric videos and text. *arXiv preprint arXiv:2210.14395*, 2022.

[28] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *European conference on computer vision*, pages 529–544. Springer, 2022.

[29] Girish Narayanswamy, Xin Liu, Kumar Ayush, Yuzhe Yang, Xuhai Xu, Shun Liao, Jake Garrison, Shyam Tailor, Jake Sunshine, Yun Liu, et al. Scaling wearable foundation models. *arXiv preprint arXiv:2410.13638*, 2024.

[30] Sameer Neupane, Mithun Saha, Nasir Ali, Timothy Hnat, Shahin Alan Samiei, Anandatirtha Nandugudi, David M Almeida, and Santosh Kumar. Momentary stressor logging and reflective visualizations: Implications for stress management with wearables. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2024.

[31] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

[32] Arvind Pillai, Dimitris Spathis, Fahim Kawsar, and Mohammad Malekzadeh. Papagei: Open foundation models for optical physiological signals. *arXiv preprint arXiv:2410.20542*, 2024.

[33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.

[34] James M Rehg, Susan A Murphy, and Santosh Kumar. *Mobile Health: Sensors, Analytic Methods, and Applications*. Springer, 2017.

[35] Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. Deep ppg: Large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19(14):3079, 2019.

[36] Aaqib Saeed, Tanir Ozcelebi, and Johan Lukkien. Multi-task self-supervised learning for human activity detection. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(2):1–30, 2019.

[37] Mithun Saha, Maxwell A Xu, Wanting Mao, Sameer Neupane, James M Rehg, and Santosh Kumar. Pulse-ppg: An open-source field-trained ppg foundation model for wearable applications across lab and field settings. *arXiv preprint arXiv:2502.01108*, 2025.

[38] 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.

[39] Mert Sevil, Mudassir Rashid, Mohammad Reza Askari, Zacharie Maloney, Iman Hajizadeh, and Ali Cinar. Detection and characterization of physical activity and psychological stress from wristband data. *Signals*, 1(2):188–208, 2020.

[40] Feng-Tso Sun, Cynthia Kuo, Heng-Tze Cheng, Senaka Buthpitiya, Patricia Collins, and Martin Griss. Activity-aware mental stress detection using physiological sensors. In *Mobile Computing, Applications, and Services: Second International ICST Conference, MobiCASE 2010, Santa Clara, CA, USA, October 25-28, 2010, Revised Selected Papers 2*, pages 282–301. Springer, 2012.

[41] Timo Sztyler and Heiner Stuckenschmidt. On-body localization of wearable devices: An investigation of position-aware activity recognition. In *2016 IEEE international conference on pervasive computing and communications (PerCom)*, pages 1–9. IEEE, 2016.

[42] Chi Ian Tang, Ignacio Perez-Pozuelo, Dimitris Spathis, and Cecilia Mascolo. Exploring contrastive learning in human activity recognition for healthcare. *arXiv preprint arXiv:2011.11542*, 2020.

[43] Rahul Thapa, Bryan He, Magnus Ruud Kjaer, Hyatt Moore IV, Gauri Ganjoo, Emmanuel Mignot, and James Y. Zou. SleepFM: Multi-modal representation learning for sleep across ECG, EEG and respiratory signals. In *AAAI 2024 Spring Symposium on Clinical Foundation Models*, 2024.

[44] Kobiljon Toshnazarov, Uichin Lee, Byung Hyung Kim, Varun Mishra, Lismer Andres Caceres Najarro, and Youngtae Noh. Sosw: Stress sensing with off-the-shelf smartwatches in the wild. *IEEE Internet of Things Journal*, 2024.

[45] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.

[46] Chongyang Wang, Yuan Feng, Lingxiao Zhong, Siyi Zhu, Chi Zhang, Siqi Zheng, Chen Liang, Yuntao Wang, Chengqi He, Chun Yu, et al. Ubiphysio: Support daily functioning, fitness, and rehabilitation with action understanding and feedback in natural language. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(1):1–27, 2024.

[47] Min Wu, Hong Cao, Hai-Long Nguyen, Karl Surmacz, and Caroline Hargrove. Modeling perceived stress via hrv and accelerometer sensor streams. In *2015 37th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pages 1625–1628. IEEE, 2015.

[48] Kang Xia, Wenzhong Li, Shiwei Gan, and Sanglu Lu. Ts2act: Few-shot human activity sensing with cross-modal co-learning. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 7(4):1–22, 2024.

[49] Maxwell Xu, Alexander Moreno, Hui Wei, Benjamin Marlin, and James Matthew Rehg. Rebar: Retrieval-based reconstruction for time-series contrastive learning. In *The Twelfth International Conference on Learning Representations*, 2024.

[50] Maxwell A Xu, Jaya Narain, Gregory Darnell, Haraldur T Hallgrimsson, Hyewon Jeong, Darren Forde, Richard Andres Fineman, Karthik Jayaraman Raghuram, James Matthew Rehg, and Shirley You Ren. Relcon: Relative contrastive learning for a motion foundation model for wearable data. In *The Thirteenth International Conference on Learning Representations*, 2024.

[51] Maxwell A Xu, Girish Narayanswamy, Kumar Ayush, Dimitris Spathis, Shun Liao, Shyam A Tailor, Ahmed Metwally, A Ali Heydari, Yuwei Zhang, Jake Garrison, et al. Lsm-2: Learning from incomplete wearable sensor data. *arXiv preprint arXiv:2506.05321*, 2025.

[52] Hang Yuan, Shing Chan, Andrew P Creagh, Catherine Tong, Aidan Acquah, David A Clifton, and Aiden Doherty. Self-supervised learning for human activity recognition using 700,000 person-days of wearable data. *NPJ digital medicine*, 7(1):91, 2024.

[53] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14730–14740, 2023.

[54] Yuwei Zhang, Kumar Ayush, Siyuan Qiao, A Ali Heydari, Girish Narayanswamy, Maxwell A Xu, Ahmed A Metwally, Shawn Xu, Jake Garrison, Xuhai Xu, et al. Sensorlm: Learning the language of wearable sensors. *arXiv preprint arXiv:2506.09108*, 2025.

A Appendix

A.1 Limitations

While ProtoMM effectively aligns multimodal biosignals, it relies on the assumption that physiological dynamics can be discretized into a finite set of shared prototypes. This design choice imposes a specific inductive bias that works well for quasi-periodic signals (e.g., cardiac cycles in PPG, gait in accelerometry) but may be suboptimal for highly stochastic or non-stationary signals lacking recurring motifs. Additionally, our shared prototype constraint assumes a symmetric semantic overlap between modalities. In scenarios where one modality contains significant unique information irrelevant to the other (e.g., motion artifacts in accelerometry unrelated to heart rate), the strict alignment objective may suppress these modality-specific features if the prototype size is not sufficiently large to accommodate them.

A.2 Reproducibility

Our Methods section in Section 3 presents the model and training setup, and Experiments section in Section 4 describes the evaluation protocol and study design. Hyperparameters for every benchmark appear in Appendix A.4. We will release our model weights and a codebase with the full training methodology, architecture, and reproducible evaluation code, upon acceptance. The PPG-DaLiA and WESAD datasets are publicly available, and Section 4.3 and Appendix A.5 explain how we curate and preprocess each one for our tasks.

A.3 Ethics

Our paper develops models using physiological signals, with the goal of improving personal health. We acknowledge the associated risks, including privacy issues and the possibility of widening health disparities, as these models enable more detailed characterization of patients. Without effective regulation, patients may have limited control over their data, raising concerns about the upholding of autonomy, a core principle of medical ethics. Our study uses de-identified data from IRB-approved protocols, ensuring no participant identification information is included in our analysis. Nevertheless, we believe our work contributes positively to the field by advancing personalized health recommendations, which can enhance care quality, patient safety, and overall well-being.

A.4 Benchmark Implementations

Our code will be made publicly available upon acceptance. All benchmark implementations adhere to the experimental setup described in Section 4, utilizing identical encoders: a 1D ResNet-26 with a kernel size of 11, a stride of 2 with global max temporal pooling to produce 512-dimensional embeddings. The only architectural difference is the number of input channels for the initial convolution layer: 3 for accelerometer data, 1 for PPG, and 4 for the unimodal methods that concatenate inputs. Each training sample is transformed into two stochastic views, and augmentations are sampled with equal probability. The set of augmentations is consistent: for the accelerometer, we use an empirically-established set comprising additive Gaussian noise, scaling, 3D rotation, negation, time reversal, channel shuffle, segment shuffle, and time warping [42]. For the single-channel PPG signal, we use the same set, omitting the inapplicable rotation and channel shuffling augmentations. Training hyperparameters remain consistent across all methods as described in Section 4.1. Below we detail the implementation-specific parameters for each baseline.

- **CLIP** [33, 43]: Pairwise contrastive alignment between temporally paired modalities. We employ modality-specific single-layer projection heads with same input and output dimensions (512). The temperature parameter is initialized to 1.0.
- **COCOA** [11]: We reimplement the original TensorFlow code (`https://github.com/cruiseresearchgroup/COCOA`) in PyTorch while maintaining architectural fidelity. Key parameters include temperature=0.1, scale_loss=1/32, and lambd=3.9e-3, window = 100. Modality-specific projectors consist of a flattening layer followed by a linear projection to maintain dimensional consistency during training. Due to the method's specific design, global pooling is disabled during training.

- **CroSSL** [10]: We reimplement the original TensorFlow code (`https://github.com/Nokia-Bell-Labs/CroSSL`) in PyTorch, which follows the original spatial masking approach with coverage=0.9. Embeddings from modality-specific encoders are processed through a shared projector network consists of several linear layers with ReLU activations in between, resulting in a final output dimension (proj_size) of 32.
- **FOCAL** [21]: We adapt the official implementation (`https://github.com/tomoyoshki/focal`) to our codebase while preserving the core methodology. Modality-specific projectors consist of two linear layers with ReLU activation. Key parameters include temperature=0.5, sequence_length=4, and loss weights: shared contrastive=1, private contrastive=1, orthogonal=3, rank=5. Augmentations include standard temporal transformations followed by frequency-domain phase shifting.
- **SLIP** [28]: We implement both within- and between-modality contrastive losses using the standard CLIP formulation with temperature initialized to 1.0.
- **SimCLR** [7]: Unimodal contrastive learning with two augmented views per sample. temperature=0.1 is used for the contrastive loss.
- **RelCon** [50]: We use the official RelCon implementation (`https://github.com/maxxu05/relcon`) and train its motif-based distance module on our pre-training datasets (MOODS, MOD). Following the default setup of the authors, we retain key hyperparameters `withinuser_cands=20` and `tau=0.1`.
- **Abbaspourazad et al.** [1]: We implement the positive pair selection and momentum training as directed in the paper with tau=0.04, momentum_tau=0.99, and koleo_lambda=0.1.

All implementations maintain fairness through consistent encoder architecture, data preprocessing, and evaluation protocols. Differences arise only in method-specific components as detailed above, ensuring meaningful comparison while preserving each approach's distinctive characteristics.

A.5 Additional Data Description

- **MOODS** contains PPG and accelerometer signals collected using a Fossil Sport (Version 4) smartwatch from 122 participants. For stress detection (Stressed vs. Unstressed), we split the dataset into 1-minute windows following previous works [26, 44]; whereas for binary activity recognition (Stationary vs. Non-stationary), we use 20-second samples for downstream evaluations.
- **WESAD** uses a wrist-worn Empatica E4 [22] to record accelerometer data (at 32 Hz) and blood volume pulse (at 64 Hz) from 15 participants. The stress detection task includes four sessions: Stress, Baseline, Amusement, and Meditation. For the binary classification task, following prior work [38, 9, 20], we drop Meditation, merge Baseline and Amusement into Non-stress, while retaining the original Stress sessions. Each session is segmented into non-overlapping 1-minute windows for downstream evaluations.
- **PPG-DaLiA** contains accelerometer data (at 32 Hz) and blood volume pulse (at 64 Hz), collected using an Empatica E4 [22] worn on the wrist and ECG (700 Hz) from chest-worn RespiBAN [4] to offer the ground truth of heart rate prediction. Fifteen subjects followed a semi-structured daily life protocol comprising eight distinct activities (Sitting, Ascending and Descending stairs, Table soccer, Cycling, Driving, Lunch break, Walking, and Working), with transient segments between activities annotated as an additional class. For model input, we segment all signals into 8-second windows with a 2-second sliding step.

A.6 Generalization to New Modalities

To demonstrate the generality of ProtoMM beyond pulse-motion data, we conduct additional experiments on two multimodal time-series benchmarks: Moving Object Detection (MOD) [21] and RealWorld-HAR [41]. These datasets allow us to evaluate two complementary forms of generalization. MOD tests whether ProtoMM transfers to non-physiological sensing modalities, while RealWorld-HAR tests whether the shared prototype objective scales beyond two-modality pairs.

MOD. MOD contains synchronized seismic vibration and acoustic recordings generated by different types of vehicles. We resample both modalities to 100 Hz and use fixed 2-second segments. We separate the data used for self-supervised pretraining and downstream evaluation following the original

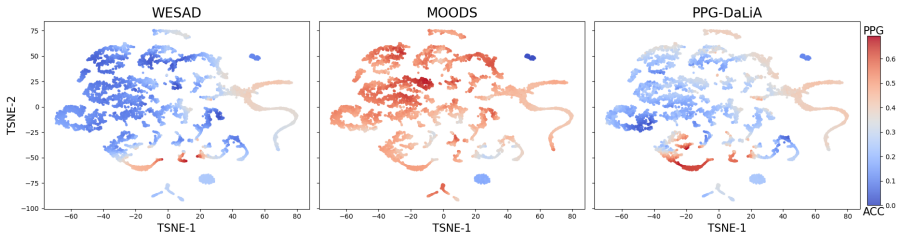

Figure 3: Across WESAD, MOODS, and PPG-DaLiA, t-SNE visualization of the learned prototype space, colored by the modality preference score. Only a smaller number of prototypes form consistent red (lower middle) or blue (upper right) subclusters across datasets, corresponding to modality-specific prototypes. The majority of prototypes exhibit varying colors across datasets, indicating that most prototypes capture shared latent states. This demonstrates that ProtoMM learns a hybrid representation consisting of mostly shared prototypes, while automatically discovering prototypes that encode complementary modality-specific patterns.

benchmark protocol in FOCAL [21]. For the pretraining split, we use a 70:15:15 train/validation/test split, where the test split is not used during pretraining. For the downstream task, we use 8:1:1 train/validation/test split. The downstream task is vehicle classification with 7 classes.

RealWorld-HAR. RealWorld-HAR contains four synchronized wearable sensing modalities: accelerometer, gyroscope, magnetometer, and light, all sampled at 50 Hz. The downstream task is activity recognition with 8 activity classes. Following FOCAL [21], we use 5-second windows with a 1-second overlap. Since RealWorld-HAR includes signals collected from multiple body positions, we combine samples across body positions to construct a unified multimodal benchmark. We exclude the forearm position because the light modality is unavailable for that position.

We perform subject-level train/validation/test splitting to avoid leakage across subjects. We hold out 20% of subjects for testing and use 20% of the remaining subjects for validation. The pretraining stage uses only the training and validation subjects. The held-out test subjects are used only for downstream evaluation.

Model architecture and training. To ensure fair comparison, all methods and modalities use the same encoder architecture: a 1D ResNet-26 with kernel size 11, stride 2, and a final embedding dimension of 128. The only architecture-dependent difference is the number of input channels for each modality-specific encoder. For RealWorld-HAR, the accelerometer, gyroscope, and magnetometer encoders each take 3-channel inputs, while the light encoder takes a single-channel input. For MOD, the acoustic and seismic encoders are treated as separate modality-specific encoders.

We use a consistent set of augmentations across methods. For inertial modalities such as accelerometer, gyroscope, and magnetometer, we use additive Gaussian noise, scaling, 3D rotation, negation, time reversal, channel shuffle, segment shuffle, and time warping [42, 50]. For single-channel modalities such as acoustic, seismic, and light signals, we use the same augmentation family while omitting transformations that are not applicable, such as 3D rotation and channel shuffling.

We optimize all models with Adam [19], using a learning rate of 10^{-5}, no weight decay, and a batch size of 256. Models are trained for 100 epochs or 9 hours, whichever comes first, on an NVIDIA L40S GPU. Training and validation losses are computed per batch and aggregated at the epoch level. For all downstream experiments, we select the checkpoint with the lowest validation loss.

Downstream evaluation. After pretraining, we freeze all encoders and evaluate the learned representations using two lightweight downstream heads: a linear probe and an MLP classifier. The MLP head is a two-layer classifier with a 64-dimensional hidden layer and ReLU activation in between. We concatenate the frozen embeddings from all modality-specific encoders before passing them to the downstream head. We report macro F1-score and accuracy for both datasets.

Table 5: Effect of the number of prototypes (P) on downstream performance.

Dataset	Task	Metric	Number of Prototypes (P)					
			0.3k	1k	3k	10k (ProtoMM)	30k	100k
MOODS	Stress (2)	↑F1	0.448±0.005	0.532±0.006	0.532±0.006	0.532±0.006	0.529±0.006	**0.543±0.006**
		↑Acc	**0.616±0.005**	0.587±0.006	0.593±0.006	0.591±0.006	0.586±0.005	0.597±0.006
	Activity (2)	↑F1	0.619±0.005	0.871±0.003	0.868±0.003	**0.872±0.003**	0.871±0.003	0.869±0.003
		↑Acc	0.855±0.002	0.931±0.002	0.929±0.002	**0.932±0.002**	0.931±0.002	0.930±0.002
WESAD	Stress (4)	↑F1	0.518±0.059	**0.641±0.053**	0.556±0.036	0.622±0.050	0.546±0.049	0.607±0.053
		↑Acc	0.596±0.055	**0.719±0.047**	0.697±0.050	**0.719±0.048**	0.618±0.052	0.697±0.049
	Stress (2)	↑F1	0.712±0.059	0.910±0.035	0.905±0.037	0.910±0.035	**0.922±0.034**	0.824±0.048
		↑Acc	0.767±0.049	0.918±0.032	0.918±0.032	0.918±0.031	**0.932±0.030**	0.836±0.044
PPG-DaLiA	Activity (9)	↑F1	0.498±0.005	**0.661±0.004**	0.656±0.004	0.656±0.004	0.645±0.004	**0.661±0.004**
		↑Acc	0.503±0.004	**0.645±0.004**	0.634±0.004	0.638±0.004	0.626±0.004	0.641±0.004
	HR (R)	↓MAE	10.58±0.08	**8.59±0.07**	8.84±0.07	8.74±0.07	8.82±0.07	8.78±0.07
		↑R^2	0.529±0.009	**0.662±0.007**	0.641±0.008	0.648±0.008	0.638±0.008	0.643±0.008

A.7 Existence of modality-specific and modality-shared prototypes.

To characterize how learned prototypes are shared or unique across modalities, we compute a *modality preference score*. For each prototype vector $\mathbf{p}_k \in \mathbb{R}^E$, we evaluate its average assignment probability under each modality samples separately. For a given dataset, this probability assignment is calculated as,

$$\mathbf{p}_k^m = \mathbb{E}_i[\mathbf{U}_{m,i}[k]] \in \mathbb{R},$$

where the expectation is taken over all samples i in modality m. We define the modality preference score as the difference between these assignment probabilities. For a PPG and Accelerometry dataset, the score of prototype $\mathbf{p}_k$ is shown as:

$$d_k = \mathbf{p}_k^{\text{PPG}} - \mathbf{p}_k^{\text{ACC}}, s_k = \frac{d_k - \min_j d_j}{\max_j d_j - \min_j d_j}, s_k \in [0, 1],$$

where $s_k \approx 1$ indicates PPG-specific prototypes and $s_k \approx 0$ indicates ACC-specific prototypes.

Figure 3 shows that only a small fraction of prototypes exhibit strong modality preference across datasets, capturing modality-specific signal characteristics. In contrast, most prototypes are activated by varying modalities across datasets, indicating that the learned representation is largely modality-invariant and reflects shared physiological structure. This explicit prototype space both blends information between modalities and preserves complementary features.

A.8 Ablation Study

A.8.1 Effect of number of prototypes

Set up. We conduct ablation study with number of prototypes P = [0.3k, 1k, 3k, 10k, 30k, 100k].

Analysis. In Table 5, across datasets, we observe stable performance when P is greater or equal to 1k. The performance gains for extremely large sets of prototypes (30k, 100k) are small (e.g. 1.1% improvement on MOODS stress F1, 1.2% improvement on WESAD binary stress F1). This confirms that ProtoMM does not rely on large prototype sets to function well. Overall, this demonstrates that ProtoMM is highly robust to prototype-set size.

A.8.2 Effect of Sinkhorn-Knopp normalization

Setup. We investigate two ablations on the Sinkhorn-Knopp normalization in ProtoMM: ProtoMM **Sinkhorn** → **Softmax** which replaces Sinkhorn normalization with Softmax, and ProtoMM **w/o Sinkhorn** that disables Sinkhorn by setting number of iterations as 0.

Analysis. Table 6 presents the results of the ablation study. Both modifications yield large performance drops across datasets (e.g. PPG-DaLiA Activity F1 drops by 30.1% and 26.9%, respectively). This highlights that Sinkhorn-Knopp is a critical mechanism for avoiding collapse and ensuring stable alignment.

Table 6: Ablation study on Sinkhorn-Knopp normalization.

	MOODS				WESAD				PPG-DaLiA			
	Stress (2)		Activity (2)		Stress (4)		Stress (2)		Activity (9)		HR (R)	
Setting	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↑F1	↑Acc	↓MAE	↑ R^2
w/ Sinkhorn (ProtoMM)	**0.532±0.006**	**0.591±0.006**	**0.872±0.003**	**0.932±0.002**	**0.622±0.050**	**0.719±0.048**	**0.910±0.035**	**0.918±0.031**	**0.656±0.004**	**0.638±0.004**	**8.74±0.07**	**0.648±0.008**
Sinkhorn → Softmax	-11.0%	+2.6%	-24.2%	-7.8%	-9.2%	-7.8%	-14.6%	-11.9%	-30.1%	-26.3%	+28.0%	-24.9%
w/o Sinkhorn	-11.6%	+1.1%	-11.9%	-4.5%	-17.0%	-15.6%	-21.8%	-16.4%	-26.9%	-22.4%	+25.6%	-23.6%

Table 7: Prototype class purity over binary classification tasks.

			Top-k Prototype Purity (%)		
Dataset	Task	Mod	Top-3	Top-5	Top-10
MOODS	Stress (2)	A	41.95	23.92	9.21
		P	36.56	16.33	4.56
	Activity (2)	A	83.57	77.65	70.13
		P	79.80	71.16	61.15
WESAD	Stress (2)	A	53.74	33.12	13.26
		P	64.85	45.62	27.38

KEY: (Mod)ality; (A)ccel, (P)PG

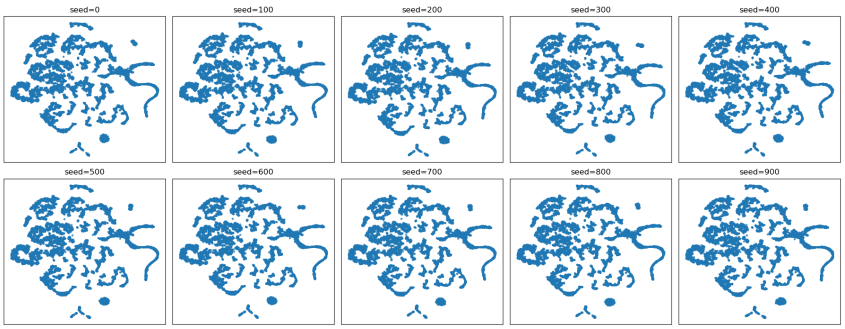

Figure 4: t-SNE visualization of prototype space across 10 different random seeds (0–900). The scatter patterns remain nearly identical, confirming that the learned prototypes form a stable structure independent of the initialization.

A.9 Prototype Class Purity

To quantify the semantic coherence of the learned prototypes, we measure the class purity of each prototype's nearest neighbors in Table 7. High purity indicates that prototype-associated embeddings cluster meaningfully according to downstream labels, despite the unsupervised training.

Across MOODS and WESAD, we observe strong purity scores, especially for activity recognition tasks. Even for stress detection, where physiological patterns are more subtle, prototypes demonstrate clear clustering of semantically aligned segments.

These results confirm that prototypes function as discrete semantic anchors in the embedding space, organizing physiological sequences into meaningful structures.

A.10 Prototypes' Stability across Initializations

We assess whether the learned prototypes depend heavily on initialization by visualizing prototype distributions across 10 random seeds. The resulting scatter plots, Figure 4, show nearly identical cluster shapes and spatial arrangements, demonstrating that prototype learning is stable and reproducible.

This robustness emerges because the prototypes act as cluster centers that converge toward meaningful physiological motifs regardless of initialization, similar to behavior in k-means clustering. The stability further reinforces that prototype-based encoder learns a genuine latent structure, rather than artifacts of random initialization.

A.11 Derivation of our Multimodal Prototype Prediction Loss from Log-likelihood

Our objective is to find the network parameters θ that maximizes the log-likelihood function of the observed i samples, composed of M modalities:

$$\theta^* = \arg\max_{\theta} \sum_i \sum_{a=1}^{A} \sum_{m=1}^{M} \log p(\mathbf{X}_{i,m}^{(a)}; \theta)$$

We assume that the observed data modality $\mathbf{X}_{i,m}^{(a)}$ is related to some combination of modality-specific and modality-shared prototypes $\mathbf{P} = \{\mathbf{P}_m \cup \mathbf{P}_s\}$ where $\mathbf{p}_{i,m}^{(a)}$ is the optimal prototype assignment of $\mathbf{X}_{i,m}^{(a)}$. We abbreviate $\mathbf{p}_{i,m}^{(a)}$ to $\mathbf{p}$ for brevity.

$$= \arg\max_{\theta} \sum_i \sum_{a=1}^{A} \sum_{m=1}^{M} \log \sum_{\mathbf{p} \in \{\mathbf{P}_m \cup \mathbf{P}_s\}} p(\mathbf{X}_{i,m}^{(a)}, \mathbf{p}; \theta)$$

Given a $Q(\mathbf{p})$ probability simplex for $\mathbf{p}$, such that $\sum_{\mathbf{p} \in \mathbf{P}} Q(\mathbf{p}) = 1$,

$$= \arg\max_{\theta} \sum_i \sum_{a=1}^{A} \sum_{m=1}^{M} \log \sum_{\mathbf{p} \in \{\mathbf{P}_m \cup \mathbf{P}_s\}} Q(\mathbf{p}) \frac{p(\mathbf{X}_{i,m}^{(a)}, \mathbf{p}; \theta)}{Q(\mathbf{p})}$$

With Jensen's inequality,

$$\geq \arg\max_{\theta} \sum_i \sum_{a=1}^{A} \sum_{m=1}^{M} \sum_{\mathbf{p} \in \{\mathbf{P}_m \cup \mathbf{P}_s\}} Q(\mathbf{p}) \log \frac{p(\mathbf{X}_{i,m}^{(a)}, \mathbf{p}; \theta)}{Q(\mathbf{p})}$$

Therefore, removing constants, we should maximize:

$$\sum_i \sum_{m=1}^{M} \sum_{\mathbf{p} \in \{\mathbf{P}_m \cup \mathbf{P}_s\}} Q(\mathbf{p}) \log p(\mathbf{X}_{i,m}^{(a)}, \mathbf{p}; \theta)$$
$$= \sum_i \sum_{m=1}^{M} \left(\sum_{\mathbf{p} \in \mathbf{P}_s} Q(\mathbf{p}) \log p(\mathbf{X}_{i,m}^{(a)}, \mathbf{p}; \theta) + \sum_{\mathbf{p} \in \mathbf{P}_m} Q(\mathbf{p}) \log p(\mathbf{X}_{i,m}^{(a)}, \mathbf{p}; \theta) \right)$$

Now with $\mathbf{U}_m^{(a)} = \mathrm{Softmax}(\mathbf{E}_m^{(a)} \mathbf{P} / \tau)$, the softmax forms a probability distribution over $\mathbf{X}_{i,m}^{(a)}$ and $\mathbf{p}$, so we can substitute it in.

$$= \sum_i \sum_{m=1}^{M} \sum_{a=1}^{A} \left(\sum_{\mathbf{p} \in \mathbf{P}_s} Q(\mathbf{p}) \log \mathbf{U}_m^{(a)} + \sum_{\mathbf{p} \in \mathbf{P}_m} Q(\mathbf{p}) \log \mathbf{U}_m^{(a)} \right)$$

Now with $\mathbf{V}_m^{(a)} = \mathrm{Sinkhorn}(\mathbf{E}_m^{(a)} \mathbf{P})$, the sinkhorn give us normalized soft prototype label assignments, we can substitute them in as being proportional to a true probability simplex.

$$\propto \sum_i \sum_{m,n=1}^{M} \sum_{a,b=1}^{A} \left(\sum_{n \neq m} \mathbf{V}_n^{(a)} \log \mathbf{U}_m^{(a)} + \sum_{b \neq a} \mathbf{V}_m^{(b)} \log \mathbf{U}_m^{(a)} \right) = -\mathcal{L}_{\text{MPP}} \cdot \beta$$

This is now equivalent our negative loss function multiplied by a constant β. Therefore, by minimizing our $\mathcal{L}_{\text{MPP}}$, we are maximizing this lower bound of the log-likelihood.