

# 000 001 002 003 004 005 A ROBUST PPG FOUNDATION MODEL USING MULTI- 006 MODAL PHYSIOLOGICAL SUPERVISION 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053

## ABSTRACT

Photoplethysmography (PPG), a non-invasive measure of changes in blood volume, is widely used in both wearable devices and clinical settings. Although recent work has explored PPG foundation models using large-scale intensive care unit (ICU) datasets, these efforts often assume the need for clean and high-quality signals. In contrast, we argue that the inherent noise and variability in ICU datasets can be harnessed to build more robust and generalizable representations. To address this, we propose a PPG foundation model that leverages accompanying electrocardiogram and respiratory signals in ICU datasets to select contrastive samples during pretraining. Our approach allows the model to retain and learn from noisy PPG segments, improving robustness without requiring multimodal inputs at inference. Our model, pretrained on 3x fewer subjects than existing state-of-the-art approaches, achieves performance improvements of up to 36% in classification and 42% in regression on 14 out of 15 diverse downstream tasks, including stress and heart rate prediction. Our results demonstrate that multimodal supervision can leverage clinical data to enable the development of robust, unimodal foundation models for both clinical and consumer-level data.

## 1 INTRODUCTION

Wearable devices are rapidly emerging as powerful tools to monitor physiological and behavioral signals in everyday life. These devices typically rely on embedded sensors that must meet strict design constraints: they must be small, low-power, cost-effective, and unobtrusive. However, these constraints often compromise signal quality, introducing noise and variability that significantly challenge downstream tasks. Consequently, there is a critical need for robust models that can learn effective representations from noisy signals while maintaining high accuracy and generalizability across diverse conditions. Among the various biosignals used in wearables, photoplethysmography (PPG) has gained prominence due to its simplicity, low energy consumption, and compatibility with optically-based sensing hardware. PPG measures changes in blood volume in peripheral tissue, allowing the estimation of vital signs such as heart rate and blood pressure (Elgendi et al., 2019). Unlike electrocardiography (ECG), PPG is more prone to motion artifacts and signal noise (Fine et al., 2021). Nevertheless, PPG’s ability to reflect vascular dynamics makes it a promising candidate for foundation models that aim to generalize across multiple tasks and conditions.

The advent of PaPaGei (Pillai et al., 2024), the first open-source PPG foundation model, highlights the growing interest in building general-purpose representations from wearable biosignals. PaPaGei demonstrates substantial performance gains over engineered features across multiple downstream tasks, including hypertension classification, blood pressure estimation, and heart rate prediction. These results suggest that foundation models can capture rich, transferable representations of PPG signals. However, PaPaGei depends on extensive preprocessing to extract clean morphological features for pretraining, and other existing approaches often rely on proprietary datasets (Saha et al., 2025; Abbaspourazad et al., 2023), limiting reproducibility and/or scalability. In this work, we investigate whether the reliance on curated, denoised PPG signals can be relaxed by leveraging co-recorded multimodal signals from large-scale intensive care unit (ICU) datasets. Specifically, we propose a novel PPG foundation model that uses co-recorded high-quality signals (electrocardiogram (ECG) and respiratory (RESP) data) to select contrastive PPG samples during pretraining. This allows us to learn from relatively noisy clinical PPG data, without requiring explicit denoising or morphological feature extraction. Importantly, our model requires only PPG signals at infer-

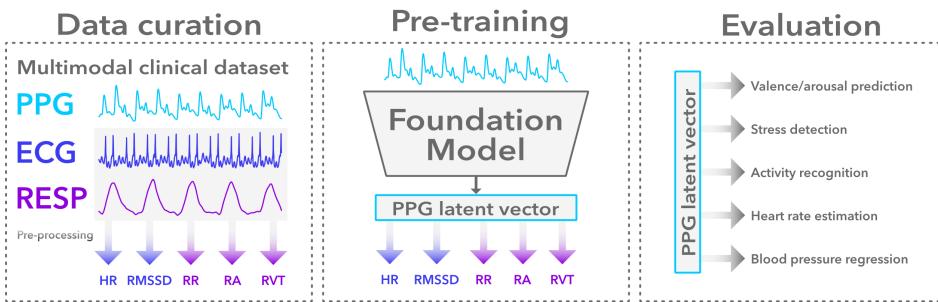


Figure 1: **Multimodal contrastive supervision framework.** (Left) The electrocardiogram (ECG) and respiratory (RESP) data co-recorded with PPG is segmented into 10s windows. Five metrics are extracted from the ECG and RESP segments that summarize those windows in a 5-dimensional vector. (Middle) The metrics are used to generative contrastive samples during pretraining. (Right) The unimodal PPG embeddings are evaluated using various tasks for unseen datasets.

ence, with multimodal data used exclusively during pretraining to enhance representation learning, as shown in Figure 1. Our key contributions are as follows:

- We demonstrate that leveraging open-source, co-recorded multimodal signals (ECG and RESP) from ICU datasets to guide PPG foundation model pretraining significantly enhances performance across a diverse set of downstream tasks.
- We introduce within-subject linear probing as a complementary evaluation method for PPG foundation models, enabling a more detailed and subject-specific assessment of representation quality and generalization beyond standard cross-subject metrics.
- Our approach outperforms PaPaGei on all but one downstream task, particularly on field-like datasets, despite using a significantly smaller subject pool and a single pretraining dataset. This highlights the efficiency, scalability, and robustness of our method.

## 2 RELATED WORK

Foundation models have demonstrated strong generalization across a wide range of domains, driven by large-scale pretraining and self-supervised learning (Bommasani, 2021). These models are often trained using self-supervised learning techniques that involve generating masked or incomplete data (Devlin et al., 2019). Generative and/or predictive pretraining has been replicated with success in other fields as well, including computer vision (He et al., 2022), pretraining for timeseries (Nie et al., 2022), and biosignals (Kostas et al., 2021; Chen et al., 2021; Chien et al., 2022; Dong et al., 2023; Liu et al., 2023; Yun et al., 2024; Zhang et al., 2024; Geenjaar & Lu, 2025). An important downside to these methods is that they can still be sensitive to noise. In practice, biosignals like PPG are far noisier than image or text data, and can be highly subject-dependent, influenced by factors such as skin tone and body composition (Bent et al., 2020). Generative approaches may thus fail to capture good embeddings for high-noise segments. Predictive approaches like JEPA (LeCun, 2022) aim to mitigate this by learning abstract representations without full reconstruction. Still, they remain sensitive to slow-varying or predictive noise patterns (Sobel et al., 2022). This slow and/or predictive noise may be induced by movement in PPG data. JEPA approaches may thus still struggle to learn good embeddings from noisy PPG data. Other, nongenerative or predictive self-supervised approaches have also been proposed for PPG data, including motif matching (Xu et al., 2023; 2024; Saha et al., 2025), morphology-based contrastive learning (Pillai et al., 2024), and temporal- or participant-based contrastive learning (Tonekaboni et al., 2021; Abbaspourazad et al., 2023). Although these approaches do not rely on generation and/or prediction during pretraining, they may still be sensitive to noise. For instance, PaPaGei relies on morphological features that are difficult to extract accurately from noisy segments, which can result in poor supervision or the exclusion of data. Moreover, similar to JEPA approaches, time- or participant-based contrastive learning that uses two noisy positive pairs may learn to focus on slow-frequency features, as opposed to potentially important medium-frequency signals like systolic peaks in the PPG signal.

108 Unimodal self-supervised supervision is only one subset of foundation models for general representation learning (Li et al., 2024). Other subsets include label supervision, multimodal supervision, and multimodal fusion models (Li et al., 2024). Given that labels vary significantly between PPG datasets, and future inference may be restricted by training a PPG foundation model using a specific set of classes, it is not practical to use label supervision for PPG foundation models. Moreover, although multimodal biosignal foundation models can be made robust to modality dropout (Liu et al., 2023; Fang et al., 2024), it is preferable to train a PPG foundation model that only requires PPG as an input to make it as general for wearable use as possible. Specifically, we focus on models that can be trained with multimodal inputs but deployed using only a single modality, making them practical for all types of wearables. In fact, large-scale open-source PPG pretraining datasets are often clinical datasets, which record multiple modalities from patients almost by default. To harness the availability of this data, we therefore focus on multimodal supervision as a method to train a PPG foundation model. Multimodal supervision models for computer vision include CLIP (Radford et al., 2021; Jia et al., 2021), and work for biosignals include multimodal contrastive pretraining (Raghu et al., 2022), BioFAME (Liu et al., 2023), SleepFM (Thapa et al., 2024), cross-modal masked auto-encoding learning (Fang et al., 2024), and CiTrus (Geenja & Lu, 2025). All of these works include various signals that are not included in most wearables, such as ECG, EOG, or EEG. Specifically for PPG data, works that use multimodal supervision for foundation models include SensorLM (Zhang et al., 2025), and work that simultaneously uses PPG and accelerometer data (Abbaspourazad et al., 2024). Both models are trained on closed-source datasets. In contrast, our approach uses open-source ICU data and relies only on PPG at inference time. We leverage biosignals co-recorded with PPG to construct a physiologically grounded supervision signal during pretraining, enabling robust representation learning from noisy clinical PPG data while maintaining scalability and reproducibility. This allows the community to build, evaluate, and extend upon our method by training on the same set of data. In fact, we share a list of the exact data files we use for our pretraining dataset in the Supplementary Material.

### 3 METHODS

136 Let  $\mathcal{D} = \{(\text{PPG}, \text{ECG}, \text{RESP})^{n,s}\}_{n=1\dots N, s=1\dots S}$  denote the multimodal biosignal dataset, where  
 137  $N$  is the number of subjects and  $S$  is the number of sessions per subject. Each tuple  
 138  $(\text{PPG}, \text{ECG}, \text{RESP})^{n,s}$  contains time-aligned signals from each modality, sampled at the same frequency and of equal length. The continuous signals are segmented into non-overlapping 10-second  
 139 windows following PaPaGei (Pillai et al., 2024),  $\mathbf{w}_t^{n,s} = (\text{PPG}, \text{ECG}, \text{RESP})_{t:(t+10s)}^{n,s}$ . The PPG segment  
 140  $\mathbf{x}_t^{n,s} = \mathbf{w}_{t,\text{PPG}}^{n,s}$  is used as the model input, while the corresponding ECG and RESP segments  
 141 are utilized to compute physiological metrics that guide contrastive supervision.

#### 3.1 A PHYSIOLOGICAL METRIC SPACE FOR CONTRASTIVE SUPERVISION

146 To ensure our PPG foundation model is robust to naturally occurring noise, such noise must be  
 147 well-represented in its pretraining dataset. Previous work relied on morphological features directly  
 148 extracted from PPG segments to define contrastive targets (Pillai et al., 2024). Noisy segments,  
 149 which are important to include in the pretraining dataset (Saha et al., 2025), consequently either  
 150 need to be discarded or lead to inaccurate contrastive targets during pretraining. By contrast, ECG  
 151 and RESP signals are typically less noisy and provide unique and physiologically relevant information.  
 152 We exploit this robustness to derive a continuous physiological metric space that reflect  
 153 the underlying cardio-respiratory state of the subject during each 10-second segment. Moreover, by  
 154 precomputing metrics from the ECG and RESP data, we can ensure metrics do not fall outside of  
 155 known physiological ranges, and filter the metrics to ensure that any noise in the ECG and RESP  
 156 data does not significantly affect our constructed physiological metric space.

157 **ECG-derived metrics.** We extract two cardiovascular targets from each 10s ECG waveform: heart  
 158 rate (**HR**) and the root mean square of successive differences (**RMSSD**). These are metrics that have  
 159 relatively good repeatability for short segments (Schroeder et al., 2004; Nussinovitch et al., 2011;  
 160 Shaffer & Ginsberg, 2017). Both provide a measure of heart rate variability (HRV), whereas resting  
 161 HR is an indicator of all-around fitness and even cardiovascular disease (Fox et al., 2007), RMSSD  
 is sensitive to autonomic function and stress (Kim et al., 2018). Especially because we are able to

162 filter the derived metrics before using them as contrastive targets, they are less sensitive to peripheral  
 163 noise than PPG and serve as a physiological target for cardiac dynamics.  
 164

165 **RESP-derived metrics.** We compute three respiratory features from each 10s RESP waveform:  
 166 respiratory rate (**RR**), respiratory amplitude (**RA**), and respiratory volume per time (**RV**). These  
 167 metrics reflect different aspects of breathing behavior, such as rhythm and tidal volume, and can  
 168 indicate stress or enhanced attention (Widjaja et al., 2013), or disorders (Brinkman et al., 2018).  
 169

170 By precomputing and filtering these metrics to ensure physiological plausibility and reduce noise-  
 171 induced artifacts, we obtain a stable, multidimensional metric space. Contrastive relationships are  
 172 then defined based on similarity in this space rather than using potentially noisy PPG morphology.  
 173 This design enables the model to learn representations that are better aligned with meaningful phys-  
 174 iological variation, and more robust to naturally occurring PPG noise.  
 175

176 **Pretraining setup and learning objective.** Given a batch of  $B$  PPG segments, two augmented  
 177 views are generated for each segment, yielding  $2B$  inputs. These are encoded using a shared con-  
 178 volutional neural network  $f_\theta(\cdot)$ , producing embeddings  $\{\mathbf{v}_i \in \mathbb{R}^{512}\}_{i=1\dots 2B}$ . Each embedding  $\mathbf{v}_i$  is  
 179 associated with a physiological metric vector  $\mathbf{y}_i$ . For each anchor embedding  $\mathbf{v}_i$ , the other embed-  
 180 dings  $\mathbf{v}_j$  ( $j \neq i$ ) are ranked according to the distance between their physiological targets  $d(\mathbf{y}_i, \mathbf{y}_j)$   
 181 in the metric space. Embeddings corresponding to more physiologically similar segments are ranked  
 182 higher. Formally, define the set  $\mathcal{S}_{i,j} = \{\mathbf{v}_k \mid k \neq i, d(\mathbf{y}_i, \mathbf{y}_k) \geq d(\mathbf{y}_i, \mathbf{y}_j)\}$ , which contains all  
 183 embeddings that are further away than  $\mathbf{v}_j$  is from  $\mathbf{v}_i$ . Then, we employ the rank-n-contrast (RNC)  
 184 loss (Zha et al., 2023), which encourages embeddings that are closer in the physiological metric  
 185 space to be closer in the learned representation space as well:  
 186

$$\mathcal{L}_{\text{RNC}} = \frac{1}{2B} \sum_{i=1}^{2B} \frac{1}{2B-1} \sum_{\substack{j=1 \\ j \neq i}}^{2B} -\log \frac{\exp(\text{sim}(\mathbf{v}_i, \mathbf{v}_j)/\tau)}{\sum_{\mathbf{v}_k \in \mathcal{S}_{i,j}} \exp(\text{sim}(\mathbf{v}_i, \mathbf{v}_k)/\tau)} \quad (1)$$

187 where  $\text{sim}(\cdot, \cdot)$  denotes cosine similarity and  $\tau$  is a temperature hyperparameter. This loss anchors  
 188 the learned PPG embeddings to the robust physiological metric space derived from multimodal  
 189 signals, improving noise robustness and encouraging physiologically meaningful representations.  
 190

## 4 EXPERIMENTAL SETUP

191 **Datasets.** For pretraining,  
 192 we use the MIMIC-III Wave-  
 193 form Database Matched Sub-  
 194 set (Goldberger et al., 2000;  
 195 Johnson et al., 2016; Moody  
 196 et al., 2020)<sup>1</sup>, which con-  
 197 tains waveform data from  
 198 10,282 ICU patients. We  
 199 selected this dataset among  
 200 the three used by PaPaGei  
 201 because its subjects are ne-  
 202 ther asleep nor under anes-  
 203 thesia, unlike the other two  
 204 datasets. This allows pa-  
 205 tients to move their arms nat-  
 206 urally, introducing realistic  
 207 movement artifacts and noise  
 208 that improve the robustness  
 209 of our model. The dataset in-  
 210 cludes multiple time-aligned  
 211 biosignals sampled at 125Hz.  
 212 The pretraining data prepro-  
 213 cessing pipeline, detailed in  
 214 Appendix A, filters the data to  
 215 retain 4,998 subjects, yield-  
 216 ing approximately 20 million 10-second PPG  
 217 segments (about 56,000 hours of data). We evaluate our model on unseen datasets and tasks, an overview of

197 Table 1: **Downstream dataset information.** More information is  
 198 provided in Appendix B. # P is short for number of participants, # S  
 199 is short for number of total samples.

Datasets	Task	Task type	# P	# S
WESAD (Schmidt et al., 2018)	Stress Affect	Clf (2) Clf (4)	15 15	4125 4125
PPG-DaLiA (Reiss et al., 2019)	Daily activities Heart rate (HR)	Clf (9) Reg	15 15	12865 64697
EEVR (Singh et al., 2024)	Valence Arousal	Clf (2) Clf (2)	37 37	10508 10508
PPG-BP (Liang et al., 2018)	Hypertension Average HR Systolic BP Diastolic BP	Clf (2) Reg Reg Reg	219 219 219 219	657 657 657 657
VitalVideos (Toye, 2023)	Systolic BP Diastolic BP	Reg Reg	100 100	300 300
WildPPG (Meier et al., 2024)	HR (green) HR (infrared) HR (red)	Reg Reg Reg	64 64 64	304708 304708 304708

216 1<sup>1</sup><https://physionet.org/content/mimic3wdb-matched/1.0/>

216 these datasets is given in Table 1. In our selection of downstream datasets we focus on wearable-  
 217 level data to verify the robustness of our model to noise in the PPG signal. Specifically, PPG-BP and  
 218 VitalVideos are similar to clinical-level PPG data, WESAD and EEVR are lab environment datasets,  
 219 and DaLiA and WildPPG are field-like datasets, which exhibit the highest noise levels.

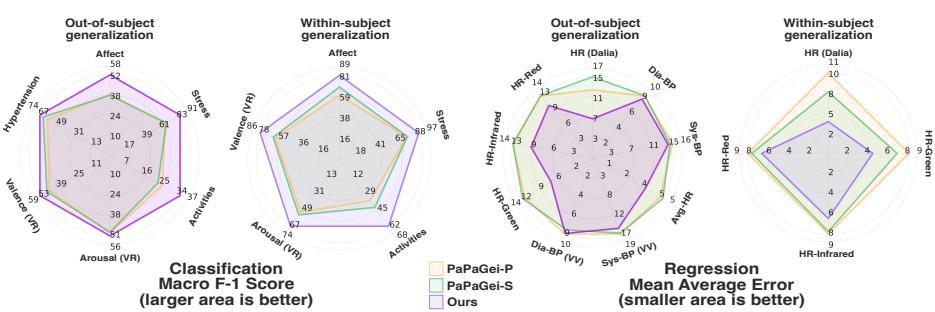
220 **Backbone & pretraining.** We use a 1D ResNet-26 convolutional encoder <sup>2</sup>  $f_\theta(\cdot)$  with instance  
 221 normalization applied to the input. The network comprises 12 residual blocks, each using a ker-  
 222 nel size of 11 and stride 2. The initial convolution outputs 128 filters, doubling every four layers.  
 223 Spatial resolution is downsampled by a factor of 2 every two layers via max-pooling. Each unfil-  
 224 tered 10-second PPG window  $\mathbf{x}^{n,t}$  is passed through the network to produce embeddings used in  
 225 our contrastive learning objective (Eq. 1). For data augmentation, two random transformations are  
 226 applied to each input window, selected from: GaussianNoise ( $p = 0.25$ ), Negation ( $p = 0.20$ ),  
 227 Scaling ( $p = 0.40$ ), and RandomCrop ( $p = 0.50$ ). These augmentations follow the same strategy  
 228 as PaPaGei (Pillai et al., 2024). Additional hyperparameter details are provided in Appendix C.

229 **Evaluation across subjects.** We evaluate our foundation model on 15 downstream tasks from six  
 230 unseen datasets, encompassing both classification and regression problems. Classification tasks  
 231 include stress, affect, arousal, valence, activity, and hypertension detection, while regression tasks  
 232 cover heart rate prediction in field, daily activity, and clinical settings, as well as diastolic and sys-  
 233 tolic blood pressure estimation. Table 1 provides a detailed list of datasets and tasks, with more  
 234 information in Appendix B.

235 Following PaPaGei (Pillai et al., 2024), we assess representation quality and generalizability us-  
 236 ing linear probing. Linear probing measures linear predictability from inferred embeddings while  
 237 keeping the backbone weights frozen. For classification tasks, we use logistic regression, and for  
 238 regression we use ridge regression. Hyperparameters for both models are tuned via 5-fold cross-  
 239 validation on the training and validation splits, and are discussed in Appendix C. Both probes are  
 240 implemented using scikit-learn (Pedregosa et al., 2011). Final results are averaged over five  
 241 test folds. Model selection uses macro F1 score for classification and mean absolute error (MAE)  
 242 for regression. Additionally, we report accuracy (ACC) and area under the receiver operating char-  
 243 acteristic curve (AUC) for classification, as well as mean squared error (MSE) and mean absolute  
 244 percentage error (MAPE) for regression.

245 **Evaluation within subjects.** To better assess model performance in realistic deployment scenarios,  
 246 we introduce a within-subject linear probing evaluation protocol. Wearable devices are typically  
 247 used by individual users, and it is often feasible to obtain labeled segments over time through user  
 248 interaction or automatic annotation. Since physiological patterns can vary significantly across in-  
 249 dividuals, evaluating linear probe performance separately for each subject provides insight into the  
 250 model’s ability to generalize under subject-specific distributions. This evaluation uses the same lin-  
 251 ear probe architecture, hyperparameter tuning strategy, and metrics as in across-subject evaluation.  
 252 However, instead of k-fold cross-validation across subjects, each fold corresponds to a user. Train-

253 <sup>2</sup><https://github.com/hsd1503/resnet1d>



256  
 257  
 258  
 259  
 260  
 261  
 262  
 263  
 264  
 265  
 266  
 267  
 268  
 269  
 Figure 2: **Comparison with state-of-the-art.** (Left) Classification results in terms of their macro  
 256 F-1 score (**larger area is better**). (Right) Regression results in terms of their mean average error  
 257 (**smaller area is better**). We evaluate across subject linear probing, and within subject linear  
 258 probing.

270 Table 2: **Downstream across-subject linear probing results.** Results are averaged over 5 test  
 271 folds; standard deviations are reported in Appendix D. For classification tasks, higher values indicate  
 272 better performance, measured by macro F1 score (MF1), accuracy (ACC), and area under the  
 273 receiver operating characteristic curve (AUC). For regression tasks, lower values are better, evaluated  
 274 using mean absolute error (MAE), mean squared error (MSE), and mean absolute percentage  
 275 error (MAPE).

277	278	279 <b>PaPaGei-P</b>			280 <b>PaPaGei-S</b>			281 <b>Ours</b>			
		282 <b>Clf (↑)</b>	283 <b>MF1</b>	284 <b>ACC</b>	285 <b>AUC</b>	286 <b>MF1</b>	287 <b>ACC</b>	288 <b>AUC</b>	289 <b>MF1</b>	290 <b>ACC</b>	291 <b>AUC</b>
292	293	Stress	0.65	0.79	0.8	0.67	0.78	0.76	0.83	0.88	0.93
294	295	Affect	0.4	0.49	0.69	0.39	0.47	0.66	0.52	0.6	0.78
296	297	Activities	0.25	0.33	0.72	0.23	0.29	0.68	0.34	0.39	0.8
298	299	Arousal	0.49	0.56	0.53	0.49	0.54	0.51	0.51	0.55	0.54
300	301	Valence	0.46	0.63	0.55	0.48	0.62	0.54	0.53	0.61	0.58
302	303	Hypertension	0.6	0.65	0.64	0.64	0.69	0.71	0.67	0.71	0.74
304	305	Avg	0.48	0.58	0.65	0.48	0.56	0.64	<b>0.57</b>	<b>0.63</b>	<b>0.73</b>
306	307	Reg (↓)	<b>MAE</b>	<b>MSE</b>	<b>MAPE</b>	<b>MAE</b>	<b>MSE</b>	<b>MAPE</b>	<b>MAE</b>	<b>MSE</b>	<b>MAPE</b>
308	309	HR (Dalia)	12.7	303	0.14	15.1	407	0.17	7.3	139	0.08
310	311	Avg-HR	4.76	40.0	0.07	4.95	43.4	0.07	3.72	23.8	0.05
312	313	Sys-BP	14.9	366	0.12	14.4	351	0.12	14.1	326	0.11
314	315	Dia-BP	8.73	121	0.12	8.77	122	0.12	8.3	114	0.12
316	317	Sys-BP (VV)	17.0	516	0.13	16.8	500	0.13	15.8	472	0.12
318	319	Dia-BP (VV)	8.7	143	0.11	8.22	125	0.1	8.63	138	0.11
320	321	HR-Green	12.2	266	0.17	12.5	273	0.17	7.41	149	0.1
322	323	HR-Infrared	12.6	273	0.17	12.7	277	0.17	9.81	212	0.14
324	325	HR-Red	12.7	279	0.17	12.7	284	0.17	10.7	233	0.15
326	327	Avg	11.6	256	0.13	11.8	265	0.14	<b>9.53</b>	<b>201</b>	<b>0.11</b>

295 ing, validation, and test sets are computed based on the temporally varying targets of that user, and  
 296 results are averaged across 5 randomly selected users, more information is provided in Appendix C.  
 297 Only datasets with temporally varying targets for each subject are included (e.g., affect detection,  
 298 heart rate), while datasets with static per-subject labels, such as PPG-BP, are excluded. This leaves  
 299 the 9 tasks from the WESAD, DaLiA, EEVR, and WildPPG datasets.

300 **Baselines.** We compare against PaPaGei-S and PaPaGei-P, the current state-of-the-art open-source  
 301 PPG foundation models. PaPaGei-S uses a morphology-based contrastive learning framework,  
 302 wherein morphological features are extracted from the raw PPG signal and used to construct positive  
 303 and negative sample pairs during pretraining (Pillai et al., 2024). In contrast, PaPaGei-P employs a  
 304 subject-aware contrastive loss, generating positive pairs from segments of the same individual and  
 305 negatives from different individuals.

## 307 5 RESULTS

309 Our model outperforms both PaPaGei-S and PaPaGei-P on nearly all (14 out of 15) downstream  
 310 tasks in across-subject and (9 out of 9) within-subject linear probing evaluations, as shown in Figure 2.  
 311 Notably, these gains are achieved despite using 3x fewer pretraining subjects than PaPaGei,  
 312 demonstrating the efficiency and robustness of our approach. Detailed numerical results for the  
 313 across-subject evaluation are provided in Table 2, while within-subject results are summarized in  
 314 Table 3. These results highlight the effectiveness and versatility of our model across diverse PPG  
 315 analysis scenarios.

316 **Evaluation across subjects.** For across-subject evaluation, our model achieves consistent and  
 317 substantial improvements over prior approaches. Notable classification gains include stress (0.83  
 318 vs. 0.67), affect detection (0.52 vs. 0.4), and daily activity classification (0.34 vs. 0.25) reflecting  
 319 stronger generalization across user states and behaviors. Regression tasks show large improvements  
 320 both under field-like (DaLiA: 7.3 vs. 12.7, WildPPG (Green): 7.41 vs. 12.2) and clinical conditions  
 321 (PPG-BP: 3.72 vs. 4.76). Even on tasks where PaPaGei specifically excels, such as systolic blood  
 322 pressure regression, our model outperforms both PaPaGei variants on the PPG-BP (14.1 vs. 14.4)  
 323 and VitalVideos (15.8 vs. 16.8) datasets. The only task our model underperforms PaPaGei-S at is  
 diastolic blood pressure regression for the VitalVideos dataset (8.63 vs. 8.22). This may be due to

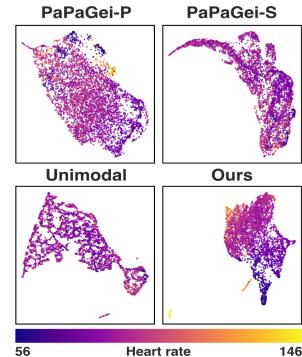
324  
325  
326  
327  
328  
329  
330 Table 3: **Within-subject downstream linear probing results.** Results are averaged over 5 test  
331 folds, with standard deviations reported in Appendix D. For classification tasks, higher values indicate  
332 better performance; metrics include macro F1 score (MF1), accuracy (ACC), and area under the  
333 ROC curve (AUC). For regression tasks, lower values are better; metrics reported are mean absolute  
334 error (MAE), mean squared error (MSE), and mean absolute percentage error (MAPE).

331	332	333	334 <b>PaPaGei-P</b>			335 <b>PaPaGei-S</b>			336 <b>Ours</b>		
			337 <b>Clf</b> ( $\uparrow$ )	338 <b>MF1</b>	339 <b>ACC</b>	340 <b>AUC</b>	341 <b>MF1</b>	342 <b>ACC</b>	343 <b>AUC</b>	344 <b>MF1</b>	345 <b>ACC</b>
346 Stress	347 0.76	348 0.83	349 0.76	350 0.75	351 0.81	352 0.77	353 0.88	354 0.91	355 0.96	356	357
358 Affect	359 0.69	360 0.73	361 0.83	362 0.61	363 0.66	364 0.83	365 0.81	366 0.84	367 0.94	368	369
370 Activities	371 0.45	372 0.51	373 0.82	374 0.39	375 0.44	376 0.75	377 0.62	378 0.62	379 0.9	380	381
382 Arousal	383 0.56	384 0.59	385 0.61	386 0.54	387 0.56	388 0.58	389 0.67	390 0.69	391 0.78	392 0.78	393 0.78
394 Valence	395 0.65	396 0.69	397 0.69	398 0.63	399 0.67	400 0.66	401 0.78	402 0.8	403 0.86	404	405
406 Avg	407 0.62	408 0.67	409 0.74	410 0.58	411 0.63	412 0.72	413 <b>0.75</b>	414 <b>0.77</b>	415 <b>0.89</b>	416	417
418 Reg ( $\downarrow$ )	419 <b>MAE</b>	420 <b>MSE</b>	421 <b>MAPE</b>	422 <b>MAE</b>	423 <b>MSE</b>	424 <b>MAPE</b>	425 <b>MAE</b>	426 <b>MSE</b>	427 <b>MAPE</b>	428	429
430 HR (Dalia)	431 7.92	432 121	433 0.09	434 10.4	435 199	436 0.12	437 4.09	438 41.2	439 0.05	440	441
442 HR-Green	443 7.22	444 122	445 0.09	446 8.45	447 186	448 0.11	449 4.67	450 73.6	451 0.06	452	453
455 HR-Infrared	456 8.03	457 142	458 0.11	459 8.27	460 151	461 0.11	462 6.68	463 118	464 0.09	465	466
468 HR-Red	469 7.85	470 130	471 0.1	472 8.26	473 185	474 0.11	475 6.82	476 111	477 0.09	478	479
482 Avg	483 7.76	484 129	485 0.1	486 8.83	487 180	488 0.11	489 <b>5.57</b>	490 <b>86.1</b>	491 <b>0.07</b>	492	493

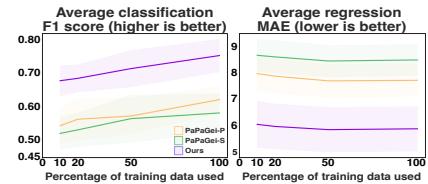
344  
345 subtle morphological differences that PaPaGei-S explicitly targets via morphology-based contrastive  
346 pretraining. In contrast, our model emphasizes robustness to noise and broader physiological varia-  
347 tion, which may trade off fine-grained waveform sensitivity in cleaner datasets. However, our model  
348 performs better on PPG-BP for the same task (8.3 vs. 8.77), and on average (8.48 vs. 8.495). Our  
349 results also replicate the general improvement of PaPaGei-S over PaPaGei-P as reported in the origi-  
350 nal PaPaGei paper (Pillai et al., 2024). Minor differences in absolute numbers in PaPaGei’s paper  
351 stem from our use of K-fold cross-validation, which better captures subject variability compared to  
352 single-split setups (Geenjaar & Lu, 2025).

353  
354 **Evaluation within subjects.** Similar to the across sub-  
355 jects results, we find that classification improvements  
356 within subjects are also higher for stress (0.88 vs. 0.76),  
357 affect (0.81 vs. 0.69), and daily activities classification  
358 (0.62 vs. 0.45). Additionally, percentage improvements  
359 relative to the best PaPaGei model for within subject  
360 valence detection are much higher than across subjects:  
361 10% across subjects, and 20% within subjects. The  
362 increased percentage improvement highlights that our  
363 model is even more accurate at tracking certain affective  
364 states in individual subjects over time. Figure 3 shows a  
365 visualization of the embedding space for a single subject  
366 from the PPG-DaLiA dataset, highlighting the difference  
367 between our method and our replication of PaPaGei-S (la-  
368 beled as ‘Unimodal’). In the figure, our model’s embed-  
369 ding space shows a clear gradient in terms of heart rate,  
370 whereas the other models do not. Moreover, since data  
371 availability for a new user may be sparse, in Figure 4  
372 we show how our model significantly outperforms both  
373 PaPaGei-P and PaPaGei-S, even on 10% data, the aver-  
374 age performance of our model is better than the best Pa-  
375 PaGei model on 100% of the data. Data is removed in a  
376 stratified manner from the training set.

377  
378 **Unimodal vs. multimodal pretraining.** To verify that  
379 differences in performance are due to our use of multi-  
380 modal contrastive guidance during pretraining, and not  
381 because of architectural and pretraining data differences,



383 Figure 3: **UMAP plots**(McInnes et al.,  
384 2018) colored by heart rate. Data are  
385 from a single subject in the PPG-DaLiA  
386 dataset.



387 Figure 4: **Average performance across**  
388 **varying percentages of within-subject**  
389 **data.** Shaded areas represent standard  
390 deviation across folds.

378 we pretrain PaPaGei-S based on the available code <sup>3</sup> on  
 379 our data. We adopt the PaPaGei-S pretraining objective,  
 380 as it was shown to consistently outperform PaPaGei-P across downstream tasks in the original work.  
 381 As a comparison, we match PaPaGei’s backbone in our model, and use our proposed multimodal  
 382 pretraining. All training hyperparameters are the same between the models, and are the same as  
 383 the ones discussed in Appendix C. As detailed in Table 4, our multimodal pretraining consistently  
 384 and substantially outperforms the unimodal PaPaGei-S baseline across all evaluated tasks, except  
 385 diastolic blood pressure regression. These results strongly validate our core hypothesis: integrating  
 386 complementary biosignal modalities during contrastive learning effectively mitigates the limitations  
 387 inherent in unimodal morphology-based contrastive targets, leading to significantly enhanced ro-  
 388 bustness, generalization, and downstream task performance.

### 389 Demographic analysis.

390 It is important to ensure  
 391 that neither the specific  
 392 set of subjects in the  
 393 pretraining dataset nor the  
 394 pretraining method lead  
 395 to demographic biases in  
 396 the model. To evaluate  
 397 how bad demographic  
 398 biases are, we use the  
 399 systolic blood pressure  
 400 regression task on the  
 401 VitalVideos dataset, which  
 402 records the Fitzpatrick skin  
 403 tone (Gupta & Sharma,  
 404 2019), age, and sex of  
 405 each subject. We use  
 406 a leave-one-subject-out  
 407 approach to perform lin-  
 408 ear probing, and also to  
 409 select hyperparameters  
 410 on the training set. The  
 411 importance of skin tone in  
 412 PPG analyses cannot be  
 413 understated because PPG  
 414 is an optical method, and  
 415 skin tone can affect light  
 416 wave reflectance (Fallow  
 417 et al., 2013). Moreover,  
 418 general device error for  
 419 wearables recording heart rate  
 420 have been found to be higher for darker skin tones (Gupta & Sharma,  
 421 2019). The results in Figure 5 show that some of the known biases appear in the models we tested.  
 422 Specifically, performance is best for the highest skin tone, and in case of skin tone 4 and 5 we see  
 423 that our model performs worse than the PaPaGei-S model. In addition, performance for adults aged  
 424 47-61 is worst and performance is lower for female subjects. In the latter case, we observe that  
 425 our model improves performance for both sexes. These findings underscore existing challenges in  
 426 equitable biosignal modeling and highlight areas for future bias mitigation.

Table 4: **Unimodal vs multimodal pre-training, same architecture and data.** Results for both methods are averaged across 5 test folds, and standard deviations can be found in Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regression tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and mean average percentage error (MAPE).

Clf ( $\uparrow$ )	Unimodal pre-training			Multimodal pre-training		
	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.63	0.79	0.84	0.76	0.83	0.86
Affect	0.39	0.51	0.71	0.43	0.53	0.71
Activities	0.19	0.32	0.7	0.31	0.37	0.78
Arousal	0.38	0.58	0.53	0.49	0.55	0.52
Valence	0.39	0.65	0.53	0.52	0.61	0.57
Hypertension	0.52	0.64	0.61	0.67	0.72	0.72
Avg	0.42	0.58	0.65	<b>0.53</b>	<b>0.6</b>	<b>0.69</b>
Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (Dalia)	16.0	443	0.18	8.25	167	0.09
Avg-HR	7.69	94.6	0.11	3.69	24.2	0.05
Sys-BP	15.8	408	0.13	14.1	326	0.11
Dia-BP	8.58	120	0.12	8.73	120	0.12
Sys-BP (VV)	16.9	505	0.13	16.2	467	0.12
Dia-BP (VV)	8.35	127	0.1	8.64	137	0.11
HR-Green	13.0	286	0.18	8.42	176	0.12
HR-Infrared	12.8	277	0.18	10.4	226	0.15
HR-Red	12.7	279	0.17	11.6	260	0.16
Avg	12.4	282	0.14	<b>10.0</b>	<b>211</b>	<b>0.11</b>

<sup>3</sup><https://github.com/Nokia-Bell-Labs/papagei-foundation-model>

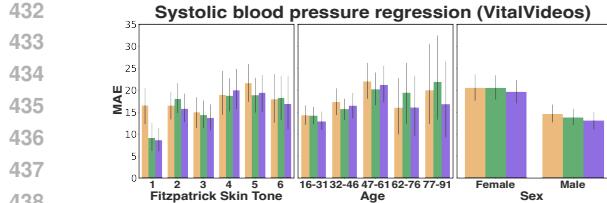


Figure 5: Systolic blood pressure regression comparison across demographic variables.

**Heart rate estimation ablations.** In Figure 6 we characterize how heart rate estimation performance varies across PPG recording locations and type of sensor (green, red, or infrared). We find that our model performs best across conditions, but especially for green wrist-worn PPG sensors. Further ablation compared with NeuroKit’s (Makowski et al., 2021) automatic heart rate estimation tool is discussed in Appendix E.

### Backbone architecture ablation.

To understand the impact of the new architecture we use, which notably has a larger number of parameters (28.8M vs. 5-5.7M), we compare the PaPaGei backbone to our proposed backbone. As shown in Table 5, our architecture yields substantial performance gains across all tasks. These results suggest that our enhanced backbone and larger model capacity contribute to improved performance, indicating potential benefits from further scaling.

## 6 CONCLUSION

This paper presents a PPG foundation model with robust multimodal pretraining, in which we use multiple biosignals alongside PPG to guide the contrastive PPG foundation model training. By using additional biosignals, we create accurate contrastive learning targets, allowing us to learn from relatively noisy clinical PPG data and improve performance on the downstream tasks. Through multiple ablation studies we demonstrate that our proposed pretraining approach greatly improves performance. In particular, our model outperforms both PaPaGei models, which are state-of-the-art PPG foundation models, in all but one out of 15 downstream tasks for across subject evaluations, and all downstream tasks for within subject evaluations. For the across subject classification and regression tasks we find improvements up to 36% for activity classification and up to 42% for field-like heart rate estimation, respectively. Notably, our model exhibits the largest performance gains for field-like datasets (DaLiA and WildPPG) and a lab dataset (WESAD). Given that these datasets focus on day-to-day PPG signals, we conclude that leveraging high-quality multimodal data during pretraining ensures our model is more robust to noise often seen in consumer-level data.



Figure 6: WildPPG heart rate estimation comparison across PPG sensor location (x-axis) and type (y-axis).

Table 5: Architecture ablation, PaPaGei backbone architecture vs our proposed architecture, same pre-training. Results for both methods are averaged across 5 test folds, and standard deviations can be found in Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regression tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and mean average percentage error (MAPE).

Clf (↑)	PaPaGei Arch			Proposed Arch		
	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.76	0.83	0.86	0.83	0.88	0.94
Affect	0.43	0.53	0.71	0.52	0.6	0.78
Activities	0.31	0.37	0.78	0.36	0.41	0.82
Arousal	0.49	0.55	0.52	0.53	0.57	0.57
Valence	0.52	0.61	0.57	0.53	0.62	0.58
Hypertension	0.67	0.72	0.72	0.71	0.75	0.77
Avg	0.53	0.6	0.69	<b>0.58</b>	<b>0.64</b>	<b>0.74</b>
Reg (↓)	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (Dalia)	8.25	167	0.09	7.78	143	0.09
Avg-HR	3.69	24.2	0.05	3.8	26.3	0.05
Sys-BP	14.1	326	0.11	13.2	281	0.11
Dia-BP	8.73	120	0.12	8.16	109	0.12
Sys-BP (VV)	16.2	467	0.12	15.9	451	0.12
Dia-BP (VV)	8.64	137	0.11	8.04	123	0.1
HR-Green	8.42	176	0.12	7.61	149	0.1
HR-Infrared	10.4	226	0.15	9.82	206	0.14
HR-Red	11.6	260	0.16	10.8	231	0.15
Avg	10.0	211	0.11	<b>9.45</b>	<b>191</b>	<b>0.11</b>

486 By using physiological metrics like heart and respiratory rate as contrastive learning targets, our  
487 method allows for a flexible combination of multiple biosignals. To further improve our current  
488 pretraining objective, one option is to include blood pressure as a target during pretraining, since it  
489 is often available in large-scale ICU datasets and is one of the tasks where our model performs most  
490 comparably to PaPaGei. Future work could also look at further improving low-data within subject  
491 performance by pretraining a model with subject-specific parameters. These parameters would then  
492 only need to be updated during a fine-tuning phase on a specific subject’s data. Another important  
493 direction for future research is to reduce the effect of skin tone on PPG foundation model embed-  
494 dings. In Figure 5 we analyze how systolic blood pressure regression performance differs across  
495 skin tones using the Fitzpatrick scale. However this scale doesn’t fully capture skin tone diversity  
496 or the biases it creates in PPG recordings (Ware et al., 2020; Colvonen, 2021). It is thus important  
497 future work to include more diverse skin tones in the ptraining data, more comprehensively test for  
498 various demographic biases, and ultimately minimize these biases.  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

540

## REPRODUCIBILITY STATEMENT

541

542 The code and weights for this paper are protected under a non-disclosure agreement (NDA) and can  
543 thus not be released to the public. To make the work as reproducible as possible, we have added  
544 a .txt file with the names of the pretraining files we use in the Supplementary Material. Moreover,  
545 both in the Methods section and Appendices **A**, **B**, and **C** we provide a very detailed explanation of  
546 our data curation, pre-processing, and experiment hyperparameters. Lastly, to make reproducing our  
547 results as easy as possible, we have used as many open-source implementations for our backbone  
548 architecture, contrastive training loss, and augmentations as possible.

548

549

## ETHICS STATEMENT

550

551 Non-invasive health monitoring can revolutionize the healthcare system. It is important to ensure  
552 that groups of people can equally benefit from non-invasive health monitoring. To verify biases in  
553 systolic blood pressure regression, we perform a demographic analysis to understand what demo-  
554 graphics our model is biased towards. Although this is a step in the right direction, it is important  
555 to develop methods that can counteract any biases and more thoroughly verify what biases exist in  
556 health foundation models. Further expansions of bias analyses across more demographics, and for  
557 more tasks is thus important, and the development of datasets that allow foundation models to be  
558 tested exhaustively before deployment is essential. The deployment of wearables for health moni-  
559 toring is also accompanied by ethical and legal implications that must be addressed (Capulli et al.,  
560 2025). Finally, it is important that any health monitoring data from wearables is used **in alignment**  
561 **with a user's preferences**. We are committed to protecting participant and/or user privacy and  
562 welfare, and to ensuring scientific validity.

562

563

## THE USE OF LARGE LANGUAGE MODELS (LLMs)

564

565 During the preparation of this manuscript, we did not substantially use LLMs. LLMs were only  
566 used to polish writing, and ensure the manuscript is approachable for a large audience. We have  
567 independently checked the correctness and clarity of the text.

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

586

587

588

589

590

591

592

593

594 REFERENCES  
595

596 Salar Abbaspourazad, Oussama Elachqar, Andrew C Miller, Saba Emrani, Udyayakumar Nallasamy,  
597 and Ian Shapiro. Large-scale training of foundation models for wearable biosignals. *arXiv  
598 preprint arXiv:2312.05409*, 2023.

599 Salar Abbaspourazad, Anshuman Mishra, Joseph Futoma, Andrew C Miller, and Ian Shapiro.  
600 Wearable accelerometer foundation models for health via knowledge distillation. *arXiv preprint  
601 arXiv:2412.11276*, 2024.

602 Brinnae Bent, Benjamin A Goldstein, Warren A Kibbe, and Jessilyn P Dunn. Investigating sources  
603 of inaccuracy in wearable optical heart rate sensors. *NPJ digital medicine*, 3(1):18, 2020.

604 Rishi Bommasani. On the opportunities and risks of foundation models. *arXiv preprint  
605 arXiv:2108.07258*, 2021.

606 Joshua E Brinkman, Fadi Toro, and Sandeep Sharma. Physiology, respiratory drive. *StatPearls*,  
607 2018.

608 Emma Capulli, Ylenia Druda, Francesco Palmese, Abdul Haleem Butt, Marco Domenicali,  
609 Anna Giulia Macchiarelli, Alessandro Silvani, Giorgio Bedogni, and Francesca Ingravallo. Eth-  
610 ical and legal implications of health monitoring wearable devices: A scoping review. *Social  
611 Science & Medicine*, pp. 117685, 2025.

612 Hugh Chen, Scott M Lundberg, Gabriel Erion, Jerry H Kim, and Su-In Lee. Forecasting adverse  
613 surgical events using self-supervised transfer learning for physiological signals. *NPJ digital  
614 medicine*, 4(1):167, 2021.

615 Hsiang-Yun Sherry Chien, Hanlin Goh, Christopher M Sandino, and Joseph Y Cheng. Maeeg:  
616 Masked auto-encoder for eeg representation learning. *arXiv preprint arXiv:2211.02625*, 2022.

617 Peter J Colvonen. Response to: investigating sources of inaccuracy in wearable optical heart rate  
618 sensors. *NPJ Digital Medicine*, 4(1):38, 2021.

619 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep  
620 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of  
621 the North American chapter of the association for computational linguistics: human language  
622 technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

623 Jiaxiang Dong, Haixu Wu, Haoran Zhang, Li Zhang, Jianmin Wang, and Mingsheng Long. Simmtm:  
624 A simple pre-training framework for masked time-series modeling. *Advances in Neural Infor-  
625 mation Processing Systems*, 36:29996–30025, 2023.

626 Mohamed Elgendi, Richard Fletcher, Yongbo Liang, Newton Howard, Nigel H Lovell, Derek Ab-  
627 bott, Kenneth Lim, and Rabab Ward. The use of photoplethysmography for assessing hyperten-  
628 sion. *NPJ digital medicine*, 2(1):60, 2019.

629 Bennett A Fallow, Takashi Tarumi, and Hirofumi Tanaka. Influence of skin type and wavelength on  
630 light wave reflectance. *Journal of clinical monitoring and computing*, 27(3):313–317, 2013.

631 Ching Fang, Christopher Sandino, Behrooz Mahasseni, Juri Minxha, Hadi Pouransari, Erdin  
632 Azemi, Ali Moin, and Ellen Zippi. Promoting cross-modal representations to improve multi-  
633 modal foundation models for physiological signals. *arXiv preprint arXiv:2410.16424*, 2024.

634 Jesse Fine, Kimberly L Branan, Andres J Rodriguez, Tananant Boonya-Ananta, Ajmal, Jessica C  
635 Ramella-Roman, Michael J McShane, and Gerard L Cote. Sources of inaccuracy in photoplethys-  
636 mography for continuous cardiovascular monitoring. *Biosensors*, 11(4):126, 2021.

637 Kim Fox, Jeffrey S Borer, A John Camm, Nicolas Danchin, Roberto Ferrari, Jose L Lopez Sendon,  
638 Philippe Gabriel Steg, Jean-Claude Tardif, Luigi Tavazzi, Michal Tendera, et al. Resting heart  
639 rate in cardiovascular disease. *Journal of the American College of Cardiology*, 50(9):823–830,  
640 2007.

648 Eloy Geenjaar and Lie Lu. Citrus: Squeezing extra performance out of low-data bio-signal transfer  
 649 learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 16781–  
 650 16789, 2025.

651

652 Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G  
 653 Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank,  
 654 physiotoolkit, and physionet: components of a new research resource for complex physiologic  
 655 signals. *circulation*, 101(23):e215–e220, 2000.

656

657 Vishal Gupta and Vinod Kumar Sharma. Skin typing: Fitzpatrick grading and others. *Clinics in*  
 658 *dermatology*, 37(5):430–436, 2019.

659

660 Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David  
 661 Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array program-  
 662 ming with numpy. *nature*, 585(7825):357–362, 2020.

663

664 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked au-  
 665 toencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*  
 666 *vision and pattern recognition*, pp. 16000–16009, 2022.

667

668 Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan  
 669 Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning  
 670 with noisy text supervision. In *International conference on machine learning*, pp. 4904–4916.  
 671 PMLR, 2021.

672

673 Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad  
 674 Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii,  
 675 a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.

676

677 Hye-Geum Kim, Eun-Jin Cheon, Dai-Seg Bai, Young Hwan Lee, and Bon-Hoon Koo. Stress and  
 678 heart rate variability: a meta-analysis and review of the literature. *Psychiatry investigation*, 15(3):  
 679 235, 2018.

680

681 Clemens Kirschbaum, Karl-Martin Pirke, and Dirk H Hellhammer. The ‘trier social stress test’–a  
 682 tool for investigating psychobiological stress responses in a laboratory setting. *Neuropsychobiol-  
 683 ogy*, 28(1-2):76–81, 1993.

684

685 Demetres Kostas, Stephane Aroca-Ouellette, and Frank Rudzicz. Bendl: Using transformers and a  
 686 contrastive self-supervised learning task to learn from massive amounts of eeg data. *Frontiers in*  
 687 *Human Neuroscience*, 15:653659, 2021.

688

689 Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open*  
 690 *Review*, 62(1):1–62, 2022.

691

692 Benjamin J Li, Jeremy N Bailenson, Adam Pines, Walter J Greenleaf, and Leanne M Williams.  
 693 A public database of immersive vr videos with corresponding ratings of arousal, valence, and  
 694 correlations between head movements and self report measures. *Frontiers in psychology*, 8:2116,  
 695 2017.

696

697 Chunyuan Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, Jianfeng Gao, et al.  
 698 Multimodal foundation models: From specialists to general-purpose assistants. *Foundations and*  
 699 *Trends® in Computer Graphics and Vision*, 16(1-2):1–214, 2024.

700

701 Yongbo Liang, Zhencheng Chen, Guiyong Liu, and Mohamed Elgendi. A new, short-recorded  
 702 photoplethysmogram dataset for blood pressure monitoring in china. *Scientific data*, 5(1):1–7,  
 703 2018.

704

705 Ran Liu, Ellen L Zippi, Hadi Pouransari, Chris Sandino, Jingping Nie, Hanlin Goh, Erdrin Azemi,  
 706 and Ali Moin. Frequency-aware masked autoencoders for multimodal pretraining on biosignals.  
 707 *arXiv preprint arXiv:2309.05927*, 2023.

702 Dominique Makowski, Tam Pham, Zen J. Lau, Jan C. Brammer, François Lespinasse, Hung Pham,  
 703 Christopher Schölzel, and S. H. Annabel Chen. NeuroKit2: A python toolbox for neurophysiological  
 704 signal processing. *Behavior Research Methods*, 53(4):1689–1696, feb 2021. doi: 10.3758/  
 705 s13428-020-01516-y. URL <https://doi.org/10.3758%2Fs13428-020-01516-y>.

706 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and  
 707 projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

708

709 Manuel Meier, Berken Utku Demirel, and Christian Holz. Wildppg: A real-world ppg dataset of  
 710 long continuous recordings. *Advances in Neural Information Processing Systems*, 37:2246–2266,  
 711 2024.

712

713 B Moody, G Moody, M Villarroel, G Clifford, and I Silva. Mimic-iii waveform database matched  
 714 subset (version1. 0). physionet, 2020.

715

716 Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64  
 717 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.

718

719 Udi Nussinovitch, Keren Politi Elishkevitz, Keren Katz, Moshe Nussinovitch, Shlomo Segev, Ben-  
 720 jamin Volovitz, and Naomi Nussinovitch. Reliability of ultra-short ecg indices for heart rate  
 721 variability. *Annals of Noninvasive Electrocardiology*, 16(2):117–122, 2011.

722

723 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, and  
 724 Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In  
 725 *Advances in Neural Information Processing Systems*, volume 32, pp. 8024–8035, 2019.

726

727 F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Pretten-  
 728 hofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and  
 729 E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*,  
 730 12:2825–2830, 2011.

731

732 Arvind Pillai, Dimitris Spathis, Fahim Kawsar, and Mohammad Malekzadeh. Papagei: Open foun-  
 733 dation models for optical physiological signals. *arXiv preprint arXiv:2410.20542*, 2024.

734

735 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
 736 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
 737 models from natural language supervision. In *International conference on machine learning*, pp.  
 738 8748–8763. PMLR, 2021.

739

740 Aniruddh Raghu, Payal Chandak, Ridwan Alam, John Guttag, and Collin Stultz. Contrastive pre-  
 741 training for multimodal medical time series. In *NeurIPS 2022 Workshop on Learning from Time  
 742 Series for Health*, 2022.

743

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

746

747 James A Russell. A circumplex model of affect. *Journal of personality and social psychology*, 39  
 748 (6):1161, 1980.

749

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

753

754 Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. In-  
 755 troducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of  
 756 the 20th ACM international conference on multimodal interaction*, pp. 400–408, 2018.

757

758 Emily B Schroeder, Eric A Whitsel, Gregory W Evans, Ronald J Prineas, Lloyd E Chambless, and  
 759 Gerardo Heiss. Repeatability of heart rate variability measures. *Journal of electrocardiology*, 37  
 760 (3):163–172, 2004.

761

762 Fred Shaffer and Jay P Ginsberg. An overview of heart rate variability metrics and norms. *Frontiers  
 763 in public health*, 5:258, 2017.

756 Pragya Singh, Ritvik Budhiraja, Ankush Gupta, Anshul Goswami, Mohan Kumar, and Pushpendra  
 757 Singh. Eevr: A dataset of paired physiological signals and textual descriptions for joint emotion  
 758 representation learning. *Advances in Neural Information Processing Systems*, 37:15765–15778,  
 759 2024.

760 Vlad Sobal, Jyothir SV, Siddhartha Jalagam, Nicolas Carion, Kyunghyun Cho, and Yann Le-  
 761 Cun. Joint embedding predictive architectures focus on slow features. *arXiv preprint*  
 762 *arXiv:2211.10831*, 2022.

763

764 Rahul Thapa, Bryan He, Magnus Ruud Kjaer, Hyatt Moore IV, Gauri Ganjoo, Emmanuel Mignot,  
 765 and James Y Zou. Sleepfm: Multi-modal representation learning for sleep across ecg, eeg and  
 766 respiratory signals. In *AAAI 2024 Spring Symposium on Clinical Foundation Models*, 2024.

767

768 Sana Tonekaboni, Danny Eytan, and Anna Goldenberg. Unsupervised representation learning for  
 769 time series with temporal neighborhood coding. *arXiv preprint arXiv:2106.00750*, 2021.

770

771 Pieter-Jan Toye. Vital videos: A dataset of face videos with ppg and blood pressure ground truths.  
 772 *arXiv preprint arXiv:2306.11891*, 2023.

773

774 Olivia R Ware, Jessica E Dawson, Michi M Shinohara, and Susan C Taylor. Racial limitations of  
 775 fitzpatrick skin type. *Cutis*, 105(2):77–80, 2020.

776

777 Devy Widjaja, Michele Orini, Elke Vlemincx, and Sabine Van Huffel. Cardiorespiratory dynamic  
 778 response to mental stress: A multivariate time-frequency analysis. *Computational and mathemat-  
 779 ical methods in medicine*, 2013(1):451857, 2013.

780

781 Maxwell A Xu, Alexander Moreno, Hui Wei, Benjamin M Marlin, and James M Rehg.  
 782 Rebar: Retrieval-based reconstruction for time-series contrastive learning. *arXiv preprint*  
 783 *arXiv:2311.00519*, 2023.

784

785 Maxwell A Xu, Jaya Narain, Gregory Darnell, Haraldur Hallgrímsson, Hyewon Jeong, Darren  
 786 Forde, Richard Fineman, Karthik J Raghuram, James M Rehg, and Shirley Ren. Relcon: Rel-  
 787 ative contrastive learning for a motion foundation model for wearable data. *arXiv preprint*  
 788 *arXiv:2411.18822*, 2024.

789

790 Taedong Yun, Justin Cosentino, Babak Behsaz, Zachary R McCaw, Davin Hill, Robert Luben, Dong-  
 791 bing Lai, John Bates, Howard Yang, Tae-Hwi Schwantes-An, et al. Unsupervised representation  
 792 learning on high-dimensional clinical data improves genomic discovery and prediction. *Nature  
 793 Genetics*, 56(8):1604–1613, 2024.

794

795 Kaiwen Zha, Peng Cao, Jeany Son, Yuzhe Yang, and Dina Katabi. Rank-n-contrast: learning con-  
 796 tinuous representations for regression. *Advances in Neural Information Processing Systems*, 36:  
 797 17882–17903, 2023.

798

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

802

803

804

805

806

807

808

809 Zexing Zhang, Huimin Lu, Songzhe Ma, Jianzhong Peng, Chenglin Lin, Niya Li, and Bingwang  
 810 Dong. A general framework for generative self-supervised learning in non-invasive estimation  
 811 of physiological parameters using photoplethysmography. *Biomedical Signal Processing and  
 812 Control*, 98:106788, 2024.

810 A ECG AND RESP PRE-PROCESSING  
811

812 We identify sessions containing more than one hour of continuous data across all three modalities:  
813 ECG, RESP, and PPG. The ECG and RESP signals are filtered using NeuroKit (Makowski et al.,  
814 2021), and then used to detect peaks: R-peaks in ECG and respiratory peaks/troughs in RESP. Signal  
815 regions without valid peaks are trimmed, and the remaining data is segmented into non-overlapping  
816 10-second windows. From the RESP signal, we compute RR, RA, and RVT, and average them  
817 within each window. From the ECG signal, we extract HR and RMSSD. Any NaN values (i.e. in  
818 case a heart rate under 30 was observed) were linearly interpolated. Afterwards, all metrics are fil-  
819 tered using a low-pass filter with a cut-off frequency of 0.001Hz (since the metrics are sampled at  
820 0.1Hz after averaging). The PPG signal is used in its raw, unfiltered form. During pretraining we  
821 map all of the metrics within the [0, 1] range to ensure that all metrics equally contribute to the dis-  
822 tance computation. We decide lower and upper bounds for the range based on known physiological  
823 ranges, and by computing the lower and upper 4 standard deviations away from the mean across all  
824 metrics in the pretraining dataset. We land on the following ranges, which we use to map all metrics  
825 between [0, 1]: HR [30, 210], RMSSD [10, 200], RA [8, 60], RR [0, 2], and RVT [0, 0.88]. Any  
826 values outside of this range are clipped to be within the range. Lastly, during pretraining we select  
827 one session for each subject for pretraining, the names of the session files that correspond to files in  
828 the [MIMIC Database](#) can be found in the Supplementary Material.  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863

864 **B DOWNSTREAM DATASET INFORMATION**  
865866 **WESAD** The Wearable Stress and Affect Detection (WESAD) dataset (Schmidt et al., 2018)  
867 contains 15 subjects recorded in a lab setting. Although the dataset records data from a variety of phys-  
868 iological sensors, we only select the PPG data, which is recorded with a 64Hz sensor. In terms of  
869 PPG preprocessing we follow (Xu et al., 2023), whose preprocessing code is available on [GitHub](#).  
870 We adapt the preprocessing code to obtain 10s non-overlapping segments, and we use PaPaGei’s  
871 `resample_batch_signal` function to resample the segments to 125Hz to match the pretraining  
872 dataset. WESAD contains 4 classes: segments that consist of a baseline recording for each subject,  
873 segments where stress is induced using a Trier Social Stress Test (TSST) (Kirschbaum et al., 1993),  
874 segments where amusement is induced using a set of eleven funny video clips, and segments where  
875 subjects follow guided meditation. In case of our affect prediction task, we classify between each  
876 of these four classes. In case of the stress classification type, we classify between segments where  
877 stress is induced versus all other segments.  
878879 **PPG-DaLiA** To better understand how well heart rate can be extracted from PPG during a wide  
880 range of activities under real-life conditions, (Reiss et al., 2019) introduced the [PPG-DaLiA dataset](#),  
881 which records PPG at 64Hz. The dataset contains data from 15 subjects that perform eight different  
882 activities: (1) Sitting still for 10 minutes (2) Ascending/descending stairs for 5 min (3) Table soccer  
883 for 5 min (4) Cycling for 8 min (5) Driving a car for 15 min (6) Having a lunch break for 30 min  
884 (7) Walking for 10 min (8) Working for 20 min. During these daily activities, both a subject’s PPG  
885 and ECG signals are recorded. The ECG signal is used as the ground truth for each 8 second win-  
886 dows, with 2 second overlap between the windows. For the heart regression task, we use this label  
887 (‘label’) because it is provided by the dataset and has been verified and preprocessed. Although  
888 the segment window is smaller than the pretraining dataset, both our and the PaPaGei backbone  
889 architecture can easily deal with slightly shorter segments because both use global averaging. For  
890 each subject, we first filter the PPG data with [PapaGei’s](#) `preprocess_one_ppg_signal`, seg-  
891 ment the data into 8 second windows with a 2 second overlap to match the target labels, and then use  
892 [PaPaGei’s](#) `resample_batch_signal` function to resample the segments to 125Hz to match the  
893 pretraining dataset. The segments are z-scored for each subject. For the daily activities classification  
894 task, we use 10 second non-overlapping segments. The labels are sampled at 4Hz (‘activity’),  
895 and we assign a label to a specific 10 second window if 75% or more of the window contains that spe-  
896 cific activity. If there is no consensus on the window, we discard it. We first filter the PPG data with  
897 [PapaGei’s](#) `preprocess_one_ppg_signal`, segment the data into 10 seconds non-overlapping  
898 segments to match the target labels, and then use [PaPaGei’s](#) `resample_batch_signal` function  
899 to resample the segments to 125Hz to match the pretraining dataset.  
900901 **EEVR** The Emotion Elicitation in Virtual Reality (EEVR) dataset (Singh et al., 2024) measures  
902 PPG data at 125Hz while 37 subjects are wearing a virtual reality (VR) headset. The study con-  
903 sists of baseline dataset collection, a VR familiarity task, and then a set of VR stimuli with post-  
904 exposure questionnaires. To evoke specific levels of arousal and valence, the authors use anno-  
905 tated 360° videos from a public database (Li et al., 2017), and select select videos based on four  
906 emotional quadrants of the Russell circumplex of affect (Russell, 1980). The circumplex con-  
907 tains two dimensions, valence and arousal, and the videos can thus be organized into high val-  
908 ence and low valence or high arousal and low arousal. The authors provide a csv file called  
909 `Raw_PPG.csv`. We use the `Participant_ID` column to separate data into specific subjects, and  
910 `Label_no_index` to separate each subject’s session into a specific video with a high/low arousal  
911 and high/low valence label. Each video’s corresponding PPG data is first filtered with [PapaGei’s](#)  
912 `preprocess_one_ppg_signal`, and data is segmented into 10 second, 5 second overlapping  
913 windows. Each segment is labeled separately for arousal and valence, and since the PPG sampling  
914 rate matches that of our pretraining dataset, we do not resample the data.  
915916 **PPG-BP** To better understand how PPG can be used to understand and predict cardiovascular  
917 disease, (Liang et al., 2018) released the PPG blood pressure dataset ([PPG-BP](#)), with PPG sampled  
918 at 1000Hz. There are three PPG recordings for each subject that last around 2 second each, and  
919 219 subjects in total. We noticed some issues with resampling the data, so we decided to linearly  
920 interpolate the data instead. Using `np.interp` (Harris et al., 2020), we interpolate each segment’s  
921 frequency from 1000Hz down to 125Hz to match the pretraining dataset frequency. Given that  
922

918 both our and PaPaGei’s architecture use global averaging, both architectures can handle a variety  
 919 of input sizes, so we didn’t pad the input data, but before interpolation we did ensure the data was  
 920 not longer than 2.1 seconds. Each subject has a recorded systolic blood pressure, diastolic blood  
 921 pressure, and hypertension label. We use the same label for each of the three segments for a subject.  
 922 After downsampling, we filter each segment with [PapaGei’s](#) `preprocess_one_ppg_signal`,  
 923 and z-score each segment.

924  
 925 **VitalVideos** As an additional evaluation of blood pressure, we also evaluate our model on the  
 926 VV-Small subset of the [VitalVideos database](#). The demographics of this dataset are outlined on  
 927 Page 6 of ([Toye, 2023](#)). The dataset contains systolic and diastolic blood pressure measurements  
 928 for 100 subjects, and PPG sampled at 55-60Hz. Given that the sampling rate varies throughout  
 929 recording, we interpolate the data to 125Hz, in order to match the sampling rate of the pretraining  
 930 dataset, with `np.interp` ([Harris et al., 2020](#)) based on the provided sample timings. We also  
 931 record the age, Fitzpatrick scale, and sex of each participant to perform our demographic analy-  
 932 sis (See Figure 5). After interpolating, the PPG data for each subject is filtered with [PapaGei’s](#)  
 933 `preprocess_one_ppg_signal`, and data is segmented into non-overlapping 10 second win-  
 934 dows. The PPG data is then z-scored for each subject. Yeah

935 **WildPPG** To better understand how different placements of PPG sensors, different types (wave-  
 936 lengths) of PPG sensors, and daily activities impact heart rate estimation ([Meier et al., 2024](#)) re-  
 937 leased the [WildPPG database](#). The dataset records data from 16 subjects, and each PPG sensor  
 938 records at 128 Hz. The ground truth estimate of the heart rate is estimated with an ECG trace  
 939 recorded from each subject’s sternum. The dataset contains data for three types of PPG sensors:  
 940 green, red, and infrared (IR), and four types of locations: wrist, head, ankle, and the sternum (chest  
 941 in our manuscript). We follow the code provided by the authors on [GitHub](#), but adapt the code in the  
 942 following ways. We ensure that the ground truth heart rate from the ECG trace is estimated in 10s  
 943 non-overlapping windows. Moreover, for each sensor, location, and subject, we filter the PPG data  
 944 with [PapaGei’s](#) `preprocess_one_ppg_signal`, segment the PPG data into non-overlapping 10s  
 945 windows, and resample the segments from 128Hz to 125Hz to match the pretraining dataset with  
 946 [PaPaGei’s](#) `resample_batch_signal`. In case it is necessary, we trim the ground-truth heart rate  
 947 segments based on the number of PPG windows. Then, we remove any segments where the ground-  
 948 truth heart rate is zero (generally indicates that the heart rate could not be estimated), and then we  
 949 z-score the PPG data for each sensor, location, and subject.

950  
 951  
 952  
 953  
 954  
 955  
 956  
 957  
 958  
 959  
 960  
 961  
 962  
 963  
 964  
 965  
 966  
 967  
 968  
 969  
 970  
 971

972 C HYPERPARAMETERS AND EXPERIMENTAL SETTINGS  
973

974 **Architecture.** Our model’s architecture is implemented in Py-  
975 Torch (Paszke et al., 2019), and consists of two main parts. First, the  
976 input to our model is a  $(2 \times \text{batch\_size}, 1, 1250)$  tensor. We use 256  
977 as the batch size during pretraining for all models. The reason we  
978 have twice as many segments along the batch dimension is because  
979 we sample two random augmentations, as described in the Methods  
980 section. This tensor first passes through an InstanceNorm1d layer,  
981 and then the 1D ResNet-26 architecture, as described in the Meth-  
982 ods section of the main text. The hyperparameters for the ResNet  
983 are shown in Table 6. The output embedding for our model is thus a  
984  $(2 \times \text{batch\_size}, 512)$  tensor. During pretraining, we attach a linear  
985 layer to the gradient-detached embeddings. The gradient detach-  
986 ing is to ensure the backbone does not train to explicitly predict the  
987 metrics during pretraining. We use the predictions from the linear  
988 layer to monitor the model’s training progress (i.e. metric predic-  
989 tions should get better during training). Moreover, after epochs 3-4  
990 we see the metric predictions get worse, which we believe indicates overfitting. Hence, we do not  
991 consider checkpoints after the first 5 epochs.

992 **Checkpoint selection.** During pretraining we save checkpoints for the backbone every 5000 steps.  
993 To select the final checkpoint that we use for comparisons, we evaluate each checkpoint on the  
994 VitalVideos Systolic BP regression task. To ensure there is not data leakage, we use each model’s  
995 training and validation set score during the hyperparameter selection process for the linear probe.  
996 The reason we use the VitalVideos Systolic BP regression task is because the results can be computed  
997 quickly. Evaluating every task for every checkpoint would require too much time and too many  
998 computational resources. In general, this metric will give us an idea about how well the model can  
999 still generalize to different datasets, and is thus valuable enough to select a checkpoint. For each  
model that we train in this paper we use a single checkpoint for all of the results.

1000 **Linear probing K-folds.** There are generally two types of datasets. Datasets where each subject  
1001 has a label. PPG-BP: average heart rate (Avg-HR), systolic blood pressure (Sys-BP), diastolic blood  
1002 pressure (Dia-BP), and hypertension, and VitalVideos: systolic blood pressure (Sys-BP VV) and  
1003 diastolic blood pressure (Dia-BP VV) from the VitalVideos dataset. The other type of dataset are  
1004 datasets with labels that vary for each subject over time. WESAD: stress and affect, PPG-DaLiA:  
1005 activities and heart rate, EEVR: arousal and valence, WildPPG: heart rate. For across subject linear  
1006 probing, datasets where each subject has a label are stratified when computing 5 folds. Specifically,  
1007 for regression tasks, values are binned into 10 bins, using an ordinal encoding, and based on quan-  
1008 tiles in the dataset using `KBinsDiscretizer`. For classification tasks, no additional binning is  
1009 necessary. The 5 splits are then obtained using `StratifiedKFold` with `random_state=42`  
1010 and shuffling on. The training indices for each fold are split into training and validation indices  
1011 with a training size of 0.75, `random_state=42`, shuffling on, and by stratifying the targets. For  
1012 datasets where labels vary over time, we obtain 5 splits with `KFold`, `random_state=42`, and  
1013 shuffling on. All names align with `scikit-learn`’s API (Pedregosa et al., 2011). For within sub-  
1014 ject linear probing, we first randomly shuffle all subjects with `np.random.default_rng(42)`,  
1015 and select the subject that corresponds to the fold index. For regression tasks, we follow the  
1016 same binning process described above to create a stratified test set. The test set is created us-  
1017 ing `train_test_split`, stratification, 0.2 as the test size, `random_state=42`, and shuffling on.  
1018 The leftover data samples are then split into a stratified training and validation set with  
1019 `train_test_split`, 0.8 as the training size, `random_state=42`, and shuffling on.

1020 **Linear probing hyperparameters.** The search space includes  $\alpha \in \{0.1, 1, 10, 100, 1000\}$  and  
1021  $\text{solver} \in \{\text{auto}, \text{cholesky}, \text{sparse\_cg}\}$ . For regression tasks, we employ ridge regression  
1022 with hyperparameters  $C \in \{0.0, 0.1, 1, 10, 100\}$  and `max_iter` fixed to 10,000. The naming of  
1023 these hyperparameters is aligned with the `scikit-learn` API (Pedregosa et al., 2011).  
1024

1025

Table 6: The 1D ResNet-26  
hyperparameters

in_channels	128
kernel_size	11
stride	2
groups	1
n_block	12
n_classes	512
downsample_gap	2
increasefilter_gap	4
use_bn	True
use_do	True
verbose	False

## 1026 D STANDARD DEVIATIONS

1028 Tables 7, 8, 9, and 10 report the standard deviations across 5 folds for the main tables in the text.  
 1029 Table 7 corresponds to Table 2 in the main text, and Table 8 corresponds to Table 3 in the main text.  
 1030 Tables 9 and 10 correspond to ablation Tables 4 and 5, respectively.

1031  
 1032 **Table 7: Downstream across subjects linear probing standard deviations across 5 folds.** Re-  
 1033 sults for each method are averaged across 5 test folds, and standard deviations can be found in  
 1034 Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1 score  
 1035 (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regres-  
 1036 sion tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and  
 1037 mean average percentage error (MAPE).

	PaPaGei-P			PaPaGei-S			Ours		
Clf ( $\uparrow$ )	MF1	ACC	AUC	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.03	0.02	0.06	0.04	0.03	0.05	0.06	0.04	0.03
Affect	0.05	0.04	0.04	0.01	0.03	0.03	0.06	0.05	0.05
Activities	0.04	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.02
Arousal	0.02	0.01	0.02	0.02	0.02	0.02	0.03	0.02	0.04
Valence	0.02	0.01	0.04	0.02	0.02	0.02	0.02	0.04	0.03
Hypertension	0.03	0.03	0.04	0.03	0.04	0.05	0.06	0.05	0.07
Avg	0.03	0.03	0.04	0.03	0.03	0.03	0.04	0.04	0.04
Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (DaLiA)	1.08	75.7	0.02	1.57	111	0.03	0.93	34.3	0.02
Avg-HR	0.27	6.04	0.0	0.35	3.56	0.0	0.24	4.56	0.0
Sys-BP	0.8	57.0	0.01	0.62	47.5	0.01	1.62	73.1	0.02
Dia-BP	0.34	6.61	0.01	0.31	6.44	0.01	0.55	11.9	0.01
Sys-BP (VV)	2.0	148	0.01	2.18	179	0.02	2.66	169	0.02
Dia-BP (VV)	1.1	56.5	0.01	1.35	53.7	0.02	1.37	64.0	0.02
HR-Green	1.51	68.1	0.04	1.28	54.4	0.03	1.36	53.7	0.03
HR-Infrared	2.05	86.1	0.04	2.17	89.5	0.05	2.08	91.0	0.04
HR-Red	0.79	29.7	0.01	0.77	27.4	0.01	0.41	11.9	0.01
Avg	1.1	59.3	0.02	1.18	63.7	0.02	1.25	57.1	0.02

1058  
 1059 **Table 8: Downstream within subjects linear probing standard deviations across 5 folds.** Re-  
 1060 sults for each method are averaged across 5 test folds, and standard deviations can be found in  
 1061 Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1 score  
 1062 (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regres-  
 1063 sion tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and  
 1064 mean average percentage error (MAPE).

	PaPaGei-P			PaPaGei-S			Ours		
Clf ( $\uparrow$ )	MF1	ACC	AUC	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.04	0.03	0.08	0.1	0.08	0.12	0.09	0.07	0.04
Affect	0.09	0.08	0.09	0.13	0.11	0.08	0.11	0.1	0.04
Activities	0.14	0.13	0.05	0.14	0.14	0.09	0.11	0.13	0.04
Arousal	0.02	0.02	0.03	0.06	0.06	0.07	0.02	0.02	0.03
Valence	0.08	0.07	0.1	0.11	0.12	0.14	0.07	0.07	0.08
Avg	0.07	0.07	0.07	0.11	0.1	0.1	0.08	0.08	0.05
Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (DaLiA)	0.88	26.7	0.01	1.03	40.0	0.01	1.02	18.9	0.02
HR-Green	0.65	36.0	0.01	0.38	48.7	0.01	1.02	36.6	0.02
HR-Infrared	1.5	69.0	0.03	1.54	80.5	0.03	1.71	71.6	0.03
HR-Red	0.44	28.6	0.01	0.5	45.8	0.01	1.06	37.4	0.02
Avg	0.87	40.1	0.02	0.86	53.7	0.02	1.2	41.1	0.02

1080

1081 **Table 9: Unimodal vs multimodal pre-training standard deviations across 5 folds.** Results for  
 1082 both methods are averaged across 5 test folds, and standard deviations can be found in Appendix D.  
 1083 For the classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1),  
 1084 accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regression  
 1085 tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and mean  
 1086 average percentage error (MAPE).

1087

	Clf ( $\uparrow$ )	Unimodal pre-training			Multimodal pre-training		
		MF1	ACC	AUC	MF1	ACC	AUC
1090	Stress	0.06	0.02	0.03	0.04	0.02	0.03
1091	Affect	0.05	0.03	0.06	0.06	0.05	0.04
1092	Activities	0.04	0.04	0.04	0.05	0.04	0.03
1093	Arousal	0.02	0.02	0.02	0.02	0.02	0.02
1094	Valence	0.0	0.0	0.02	0.01	0.03	0.02
1095	Hypertension	0.05	0.04	0.04	0.02	0.02	0.03
1096	Avg	0.04	0.02	0.03	0.03	0.03	0.03
1097	Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE
1098	HR (DaLiA)	2.64	167	0.04	1.0	36.8	0.03
1099	Avg-HR	0.42	11.4	0.01	0.27	5.86	0.0
1100	Sys-BP	0.37	33.2	0.0	0.54	50.8	0.01
1101	Dia-BP	0.23	6.69	0.01	0.53	12.1	0.01
1102	Sys-BP (VV)	1.76	142	0.01	1.45	98.9	0.01
1103	Dia-BP (VV)	1.49	61.3	0.02	1.43	67.3	0.02
1104	HR-Green	1.22	52.7	0.03	1.59	66.0	0.03
1105	HR-Infrared	2.19	89.0	0.05	2.07	88.4	0.04
1106	HR-Red	0.73	28.2	0.01	0.82	31.5	0.01
1107	Avg	1.23	65.8	0.02	1.08	50.8	0.02

1108

1109 **Table 10: Architecture ablation standard deviations across 5 folds.** Results for both methods  
 1110 are averaged across 5 test folds, and standard deviations can be found in Appendix D. For the  
 1111 classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1), accuracy  
 1112 (ACC), and the area under the receiver operating characteristic (AUC). For regression tasks lower  
 1113 is better, and we use mean average error (MAE), mean squared error (MSE), and mean average  
 1114 percentage error (MAPE).

1115

	Clf ( $\uparrow$ )	PaPaGei Arch			Proposed Arch		
		MF1	ACC	AUC	MF1	ACC	AUC
1117	Stress	0.04	0.02	0.03	0.06	0.04	0.03
1118	Affect	0.06	0.05	0.04	0.06	0.05	0.05
1119	Activities	0.05	0.04	0.03	0.03	0.03	0.02
1120	Arousal	0.02	0.02	0.02	0.03	0.02	0.04
1121	Valence	0.01	0.03	0.02	0.02	0.04	0.03
1122	Hypertension	0.02	0.02	0.03	0.06	0.05	0.07
1123	Avg	0.03	0.03	0.03	<b>0.04</b>	<b>0.04</b>	<b>0.04</b>
1124	Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE
1125	HR (DaLiA)	1.0	36.8	0.03	0.93	34.3	0.02
1126	Avg-HR	0.27	5.86	0.0	0.24	4.56	0.0
1127	Sys-BP	0.54	50.8	0.01	1.62	73.1	0.02
1128	Dia-BP	0.53	12.1	0.01	0.55	11.9	0.01
1129	Sys-BP (VV)	1.45	98.9	0.01	2.66	169	0.02
1130	Dia-BP (VV)	1.43	67.3	0.02	1.37	64.0	0.02
1131	HR-Green	1.59	66.0	0.03	1.36	53.7	0.03
1132	HR-Infrared	2.07	88.4	0.04	2.08	91.0	0.04
1133	HR-Red	0.82	31.5	0.01	0.41	11.9	0.01
1134	Avg	1.08	50.8	0.02	1.25	57.1	0.02

1134 **E HEART RATE ESTIMATION WITH NEUROKIT ABLATION**  
1135

1136 Heart rate estimation is often done using automated tools, but in cases where PPG segments are  
 1137 quite noisy, they may fail. In Table 11 we compare the best-performing PaPaGei model in terms of  
 1138 heart rate estimation (PaPaGei-P) with NeuroKit’s automatic heart rate estimation, and our proposed  
 1139 model. All models take 10s of PPG segments as input, and for NeuroKit (Makowski et al., 2021)  
 1140 we use `ppg_process`’s `PPG_Rate` output. If no heart rate was detected or not enough peaks were  
 1141 present for NeuroKit, the heart rate was set to 0. Afterwards, we perform the linear probing proce-  
 1142 dure to account for small linear errors in the NeuroKit model, and to make the procedure as similar  
 1143 to the results reported for PaPaGei-P and our model. Although the example is a little manufactured  
 1144 given that NeuroKit is often used to estimate heart rate for longer segments of PPG data, our exper-  
 1145 iment provides a one-to-one comparison for real-time 10s window heart rate estimation. Moreover,  
 1146 Table 11 shows that our model outperforms both models, and that PaPaGei outperforms NeuroKit.  
 1147 In some cases, PaPaGei-P and NeuroKit’s performances are closely matched, e.g. for HR (DaLiA)  
 1148 and HR-Green.

1149 **Table 11: NeuroKit ablation, NeuroKit heart rate estimation vs. PaPaGei-P and our model.**  
 1150 Results for both methods are averaged across 5 test folds, and standard deviations can be found  
 1151 in Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1  
 1152 score (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For  
 1153 regression tasks lower is better, and we use mean average error (MAE), mean squared error (MSE),  
 1154 and mean average percentage error (MAPE).

	NeuroKit			PaPaGei-P			Ours		
Clf ( $\uparrow$ )	MF1	ACC	AUC	MF1	ACC	AUC	MF1	ACC	AUC
Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (DaLiA)	12.9	318	0.14	12.7	303	0.14	7.3	139	0.08
Avg-HR	8.67	116	0.12	4.76	40.0	0.07	3.72	23.8	0.05
HR-Green	12.1	256	0.16	12.2	266	0.17	7.41	149	0.1
HR-Infrared	13.1	289	0.18	12.6	273	0.17	9.81	212	0.14
HR-Red	13.1	291	0.18	12.7	279	0.17	10.7	233	0.15
Avg	11.9	254	0.16	11.0	232	0.14	<b>7.79</b>	<b>151</b>	<b>0.1</b>

1164  
 1165  
 1166  
 1167  
 1168  
 1169  
 1170  
 1171  
 1172  
 1173  
 1174  
 1175  
 1176  
 1177  
 1178  
 1179  
 1180  
 1181  
 1182  
 1183  
 1184  
 1185  
 1186  
 1187

1188 F EXPANDED BASELINE RESULTS  
11891190 Table 12: **Additional baseline results.** Results for each method are averaged across 5 test folds,  
1191 and standard deviations can be found in Appendix D. For the classification tasks, higher is better,  
1192 and evaluation metrics are macro F-1 score (MF1), accuracy (ACC), and the area under the receiver  
1193 operating characteristic (AUC). For regression tasks lower is better, and we use mean average error  
1194 (MAE), mean squared error (MSE), and mean average percentage error (MAPE).  
1195

1196 1197	1198 <b>BYOL</b>			1199 <b>SimCLR</b>			1200 <b>PulsePPG</b>			
	1201 <b>Clf (↑)</b>	1202 <b>MF1</b>	1203 <b>ACC</b>	1204 <b>AUC</b>	1205 <b>MF1</b>	1206 <b>ACC</b>	1207 <b>AUC</b>	1208 <b>MF1</b>	1209 <b>ACC</b>	1210 <b>AUC</b>
1211 Stress	1212 0.81	1213 0.87	1214 0.92	1215 0.83	1216 0.88	1217 0.94	1218 0.77	1219 0.77	1220 0.84	1221 0.88
1222 Affect	1223 0.49	1224 0.56	1225 0.75	1226 0.5	1227 0.59	1228 0.77	1229 0.45	1230 0.54	1231 0.54	1232 0.73
1233 Activities	1234 0.33	1235 0.38	1236 0.8	1237 0.32	1238 0.38	1239 0.8	1240 0.36	1241 0.4	1215 0.81	1216 0.81
1215 Arousal	1216 0.51	1217 0.57	1218 0.55	1219 0.52	1220 0.56	1221 0.56	1222 0.5	1223 0.55	1224 0.54	1225 0.54
1226 Valence	1227 0.52	1228 0.61	1229 0.58	1230 0.53	1231 0.61	1232 0.58	1233 0.54	1234 0.6	1235 0.57	1236 0.57
1237 Hypertension	1238 0.69	1239 0.73	1240 0.73	1241 0.68	1215 0.72	1216 0.74	1217 0.7	1218 0.74	1219 0.75	1214 0.75
1215 Avg	1216 0.56	1217 0.62	1218 0.72	1219 0.56	1220 0.62	1221 0.73	1222 0.55	1223 0.61	1224 0.71	1225 0.71
1226 <b>Reg (↓)</b>	1227 <b>MAE</b>	1228 <b>MSE</b>	1229 <b>MAPE</b>	1230 <b>MAE</b>	1231 <b>MSE</b>	1232 <b>MAPE</b>	1233 <b>MAE</b>	1234 <b>MSE</b>	1235 <b>MAPE</b>	
	1236 HR (Dalia)	1237 9.08	1238 176	1239 0.1	1240 9.49	1241 186	1215 0.11	1216 9.03	1217 170	1218 0.1
1215 Avg-HR	1216 3.98	1217 29.3	1218 0.05	1219 4.4	1220 33.6	1221 0.06	1222 4.17	1223 30.9	1224 0.06	1225 0.06
1226 Sys-BP	1227 13.9	1228 314	1229 0.11	1230 13.6	1231 308	1232 0.11	1233 13.6	1234 313	1235 0.11	1236 0.11
1237 Dia-BP	1238 8.26	1239 110	1240 0.12	1241 8.31	1215 113	1216 0.12	1217 8.42	1218 114	1219 0.12	1220 0.12
1226 Sys-BP (VV)	1227 16.8	1228 492	1229 0.13	1230 16.0	1231 458	1232 0.12	1233 15.2	1234 400	1235 0.11	1236 0.11
1237 Dia-BP (VV)	1238 8.44	1239 125	1240 0.1	1241 8.17	1215 122	1216 0.1	1217 8.09	1218 118	1219 0.1	1220 0.1
1226 HR-Green	1227 8.67	1228 171	1229 0.12	1230 8.8	1231 174	1232 0.12	1233 9.47	1234 188	1235 0.13	1236 0.13
1237 HR-Infrared	1238 10.7	1239 228	1240 0.15	1241 10.8	1215 230	1216 0.15	1217 10.8	1218 221	1219 0.15	1220 0.15
1226 HR-Red	1227 11.4	1228 248	1229 0.16	1230 11.5	1231 250	1232 0.16	1233 12.1	1234 261	1235 0.17	1236 0.17
1237 Avg	1238 10.1	1239 210	1240 <b>0.12</b>	1241 10.1	1215 208	1216 0.12	1217 <b>10.1</b>	1218 <b>202</b>	1219 0.12	1220 0.12

1242  
 1243  
 1244  
 1245  
 1246  
 1247  
 1248  
 1249  
 1250  
 1251  
 1252  
 1253  
 1254  
 1255  
 1256

1257 **Table 13: Baselines across subjects linear probing results compared to our model.** Results for  
 1258 each method are averaged across 5 test folds, and standard deviations can be found in Appendix D.  
 1259 For the classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1),  
 1260 accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regression  
 1261 tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and mean  
 1262 average percentage error (MAPE).

	PulsePPG			SimCLR			Ours		
Clf ( $\uparrow$ )	MF1	ACC	AUC	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.77	0.84	0.88	0.83	0.88	0.94	0.83	0.88	0.94
Affect	0.45	0.54	0.73	0.5	0.59	0.77	0.52	0.6	0.78
Activities	0.36	0.4	0.81	0.32	0.38	0.8	0.36	0.41	0.82
Arousal	0.5	0.55	0.54	0.52	0.56	0.56	0.53	0.57	0.57
Valence	0.54	0.6	0.57	0.53	0.61	0.58	0.53	0.62	0.58
Hypertension	0.7	0.74	0.75	0.68	0.72	0.74	0.71	0.75	0.77
Avg	0.55	0.61	0.71	0.56	0.62	0.73	<b>0.58</b>	<b>0.64</b>	<b>0.74</b>
Reg ( $\downarrow$ )	MAE	MSE	MAPE	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (Dalia)	9.03	170	0.1	9.49	186	0.11	7.78	143	0.09
Avg-HR	4.17	30.9	0.06	4.4	33.6	0.06	3.8	26.3	0.05
Sys-BP	13.6	313	0.11	13.6	308	0.11	13.2	281	0.11
Dia-BP	8.42	114	0.12	8.31	113	0.12	8.16	109	0.12
Sys-BP (VV)	15.2	400	0.11	16.0	458	0.12	15.9	451	0.12
Dia-BP (VV)	8.09	118	0.1	8.17	122	0.1	8.04	123	0.1
HR-Green	9.47	188	0.13	8.8	174	0.12	7.61	149	0.1
HR-Infrared	10.8	221	0.15	10.8	230	0.15	9.82	206	0.14
HR-Red	12.1	261	0.17	11.5	250	0.16	10.8	231	0.15
Avg	10.1	202	0.12	10.1	208	0.12	<b>9.45</b>	<b>191</b>	<b>0.11</b>

1282  
 1283  
 1284  
 1285  
 1286  
 1287  
 1288  
 1289  
 1290  
 1291  
 1292  
 1293  
 1294  
 1295

## G MULTIMODAL METRIC ABLATION

Table 14: **ECG vs RESP ablation.** Results for each method are averaged across 5 test folds, and standard deviations can be found in Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regression tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and mean average percentage error (MAPE).

Clf (↑)	ECG			RESP			ECG + RESP		
	MF1	ACC	AUC	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.83	0.88	0.94	0.83	0.88	0.94	0.83	0.88	0.94
Affect	0.53	0.61	0.78	0.52	0.6	0.77	0.52	0.6	0.78
Activities	0.36	0.41	0.82	0.36	0.41	0.81	0.36	0.41	0.82
Arousal	0.52	0.54	0.53	0.53	0.57	0.56	0.53	0.57	0.57
Valence	0.54	0.61	0.57	0.54	0.61	0.58	0.53	0.62	0.58
Hypertension	0.7	0.73	0.76	0.7	0.73	0.75	0.71	0.75	0.77
Avg	0.58	0.63	0.73	0.58	0.63	0.74	<b>0.58</b>	<b>0.64</b>	<b>0.74</b>
Reg (↓)	MAE	MSE	MAPE	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (Dalia)	6.9	123	0.08	8.97	170	0.1	7.78	143	0.09
Avg-HR	3.66	23.2	0.05	4.06	28.0	0.06	3.8	26.3	0.05
Sys-BP	13.2	289	0.11	13.4	288	0.11	13.2	281	0.11
Dia-BP	8.38	113	0.12	8.3	110	0.12	8.16	109	0.12
Sys-BP (VV)	15.9	457	0.12	15.5	438	0.12	15.9	451	0.12
Dia-BP (VV)	8.33	130	0.1	8.11	116	0.1	8.04	123	0.1
HR-Green	7.17	142	0.1	8.53	167	0.12	7.61	149	0.1
HR-Infrared	9.38	195	0.13	10.6	222	0.15	9.82	206	0.14
HR-Red	10.4	222	0.14	11.4	246	0.16	10.8	231	0.15
Avg	<b>9.25</b>	<b>188</b>	<b>0.1</b>	9.87	198	0.11	9.45	191	0.11

Table 15: **HR vs HR + RESP ablation.** Results for each method are averaged across 5 test folds, and standard deviations can be found in Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1 score (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For regression tasks lower is better, and we use mean average error (MAE), mean squared error (MSE), and mean average percentage error (MAPE).

Clf (↑)	HR			HR + RESP			ECG + RESP		
	MF1	ACC	AUC	MF1	ACC	AUC	MF1	ACC	AUC
Stress	0.83	0.88	0.94	0.83	0.88	0.94	0.83	0.88	0.94
Affect	0.53	0.6	0.77	0.52	0.6	0.77	0.52	0.6	0.78
Activities	0.35	0.41	0.82	0.36	0.41	0.82	0.36	0.41	0.82
Arousal	0.5	0.56	0.53	0.53	0.56	0.55	0.53	0.57	0.57
Valence	0.53	0.61	0.56	0.54	0.61	0.58	0.53	0.62	0.58
Hypertension	0.68	0.7	0.73	0.71	0.74	0.76	0.71	0.75	0.77
Avg	0.57	0.63	0.73	<b>0.58</b>	0.63	0.74	0.58	<b>0.64</b>	<b>0.74</b>
Reg (↓)	MAE	MSE	MAPE	MAE	MSE	MAPE	MAE	MSE	MAPE
HR (Dalia)	6.82	123	0.08	8.29	153	0.09	7.78	143	0.09
Avg-HR	3.71	24.7	0.05	4.14	28.6	0.06	3.8	26.3	0.05
Sys-BP	13.6	309	0.11	13.3	289	0.11	13.2	281	0.11
Dia-BP	8.28	112	0.12	8.18	109	0.12	8.16	109	0.12
Sys-BP (VV)	16.2	453	0.12	15.9	454	0.12	15.9	451	0.12
Dia-BP (VV)	8.1	123	0.1	8.13	120	0.1	8.04	123	0.1
HR-Green	7.09	139	0.1	7.99	154	0.11	7.61	149	0.1
HR-Infrared	9.33	194	0.13	10.2	214	0.14	9.82	206	0.14
HR-Red	10.4	223	0.14	11.2	241	0.15	10.8	231	0.15
Avg	<b>9.29</b>	<b>189</b>	<b>0.1</b>	9.7	196	0.11	9.45	191	0.11

## 1350 H CONTRASTIVE LEARNING ABLATION

1352  
 1353 Table 16: **Our model with L2 vs the contrastive loss.** Results for each method are averaged across  
 1354 5 test folds, and standard deviations can be found in Appendix D. For the classification tasks, higher  
 1355 is better, and evaluation metrics are macro F-1 score (MF1), accuracy (ACC), and the area under  
 1356 the receiver operating characteristic (AUC). For regression tasks lower is better, and we use mean  
 1357 average error (MAE), mean squared error (MSE), and mean average percentage error (MAPE).

1358	1359	1360 <b>Ours (L2 loss)</b>			1361 <b>Ours (new)</b>		
		1362 <b>Clf (↑)</b>	1363 <b>MF1</b>	1364 <b>ACC</b>	1365 <b>AUC</b>	1366 <b>MF1</b>	1367 <b>ACC</b>
1368	1369	1370 Stress	0.81	0.86	0.92	0.83	0.88
1371	1372	1373 Affect	0.5	0.58	0.76	0.52	0.6
1374	1375	1376 Activities	0.31	0.37	0.79	0.36	0.41
1377	1378	1379 Arousal	0.51	0.56	0.54	0.53	0.57
1379	1380	1381 Valence	0.53	0.61	0.59	0.53	0.62
1381	1382	1383 Hypertension	0.69	0.73	0.73	0.71	0.75
1382	1383	1384 Avg	0.56	0.62	0.72	<b>0.58</b>	<b>0.64</b>
1383	1384	1385 Reg (↓)	<b>MAE</b>	<b>MSE</b>	<b>MAPE</b>	<b>MAE</b>	<b>MSE</b>
1384	1385	1386 HR (Dalia)	8.39	163	0.09	7.78	143
1385	1386	1387 Avg-HR	3.84	26.1	0.05	3.8	26.3
1386	1387	1388 Sys-BP	14.1	330	0.11	13.2	281
1387	1388	1389 Dia-BP	8.64	121	0.12	8.16	109
1388	1389	1390 Sys-BP (VV)	15.6	433	0.12	15.9	451
1389	1390	1391 Dia-BP (VV)	8.37	126	0.1	8.04	123
1390	1391	1392 HR-Green	8.52	173	0.12	7.61	149
1391	1392	1393 HR-Infrared	10.8	237	0.15	9.82	206
1392	1393	1394 HR-Red	11.6	255	0.16	10.8	231
1393	1394	1395 Avg	9.98	207	0.11	<b>9.45</b>	<b>191</b>
1394	1395	1396					<b>0.11</b>
1395	1396	1397					
1396	1397	1398					
1397	1398	1399					
1398	1399	1400					
1399	1400	1401					
1400	1401	1402					
1401	1402	1403					

## 1404 I PROJECTOR ABLATION

1405  
 1406 Given the success of the BYOL and SimCLR baselines, we decided to pre-train a new version of  
 1407 our model that also has a projector that we calculate the Rank-N-Contrast loss on, instead of on the  
 1408 embeddings directly. We found gains in both the classification and regression tasks, and use the  
 1409 updated model for all experiments.

1410  
 1411 **Table 17: New model (with projector) vs original model across subjects linear probing results.**  
 1412 Results for each method are averaged across 5 test folds, and standard deviations can be found  
 1413 in Appendix D. For the classification tasks, higher is better, and evaluation metrics are macro F-1  
 1414 score (MF1), accuracy (ACC), and the area under the receiver operating characteristic (AUC). For  
 1415 regression tasks lower is better, and we use mean average error (MAE), mean squared error (MSE),  
 1416 and mean average percentage error (MAPE).

1417 1418	1419 <b>Ours</b>			1420 <b>Ours (new)</b>		
	1421 <b>Clf (↑)</b>	1422 <b>MF1</b>	1423 <b>ACC</b>	1424 <b>AUC</b>	1425 <b>MF1</b>	1426 <b>ACC</b>
1427 Stress	0.82	0.87	0.93	0.83	0.88	0.94
1428 Affect	0.53	0.62	0.78	0.52	0.6	0.78
1429 Activities	0.35	0.41	0.81	0.36	0.41	0.82
1430 Arousal	0.53	0.56	0.55	0.53	0.57	0.57
1431 Valence	0.55	0.62	0.59	0.53	0.62	0.58
1432 Hypertension	0.68	0.72	0.75	0.71	0.75	0.77
1433 Avg	0.58	0.63	0.74	<b>0.58</b>	<b>0.64</b>	<b>0.74</b>
1434 1435	1436 <b>Reg (↓)</b>			1437 <b>MAE</b>		
	1438 <b>MAE</b>	1439 <b>MSE</b>	1440 <b>MAPE</b>	1441 <b>MAE</b>	1442 <b>MSE</b>	1443 <b>MAPE</b>
1444 HR (Dalia)	7.2	134	0.08	7.78	143	0.09
1445 Avg-HR	3.74	24.9	0.05	3.8	26.3	0.05
1446 Sys-BP	13.5	299	0.11	13.2	281	0.11
1447 Dia-BP	8.38	115	0.12	8.16	109	0.12
1448 Sys-BP (VV)	16.6	473	0.12	15.9	451	0.12
1449 Dia-BP (VV)	9.0	145	0.11	8.04	123	0.1
1450 HR-Green	7.28	144	0.1	7.61	149	0.1
1451 HR-Infrared	9.57	201	0.13	9.82	206	0.14
1452 HR-Red	10.6	227	0.15	10.8	231	0.15
1453 Avg	9.54	196	0.11	<b>9.45</b>	<b>191</b>	<b>0.11</b>