RAMP: RELIABILITY-ADAPTIVE MASKING FOR PRETRAINING NOISY MULTIMODAL PHYSIOLOGICAL SIGNALS

Anonymous authors
Paper under double-blind review

ABSTRACT

Physiological signals from wearable and clinical monitors are often corrupted by motion artifacts, baseline drift, and contact noise, yet most self-supervised pretraining methods apply fixed masking policies regardless of local reliability. so reconstruction can follow artifact structure instead of physiological semantics. We present RAMP, a reliability-adaptive framework that turns local signal quality into a joint pretraining policy. A Reliability Profiler describes every ECG and PPG patch through quality, corruption type, and spectral degradation. A compact Reliability Policy Network then selects the masking threshold, temporal granularity, and cross-modal fusion weights through differentiable actions, while Reliability-Weighted Reconstruction (RWR) keeps degraded but informative targets from being ignored. The resulting objective uses cleaner temporal and cross-modal context without changing the encoder used for downstream transfer. On MIMIC-III and VitalDB, RAMP reduces pretraining RMSE by 16% relative to the strongest fixed-policy baseline, with the largest benefits under asymmetric modality corruption. Across ten seeds, the pretrained encoder reaches 93.9/93.9/93.2% mean ACC/macro-F1/ROC-AUC on six benchmarks covering ECG, PPG, EEG, and accelerometry, improving over reproduced baselines on most downstream benchmarks under a unified evaluation protocol. The evidence supports reliability-aware masking as a practical way to improve representation learning when corruption is heterogeneous, while the reported scope remains controlled synthetic corruption.

1 INTRODUCTION

Continuous physiological monitoring is affected by motion artifacts, baseline drift, and contact noise whose rates change with activity, posture, and sensor contact (Liu et al., 2025; Vila et al., 2021; Scardulla et al., 2023; Gambhir et al., 2021; Krishnan et al., 2022). Masked reconstruction (He et al., 2022; Zhang et al., 2024), joint-embedding prediction (Assran et al., 2023; Thimonier et al., 2025), and contrastive pretraining (Chen et al., 2020) usually fix patch-selection or augmentation policies before observing local quality. Clean and degraded patches therefore receive the same treatment, allowing reconstruction to follow artifact-consistent structure rather than physiological semantics.

The problem is sharper for multimodal signals. ECG and PPG degrade through different physical mechanisms—electrode contact noise and baseline drift in ECG versus optical-path motion interference and ambient-light contamination in PPG—and their unreliable patches need not coincide (Park et al., 2022; Liu et al., 2024). Existing quality assessment provides useful per-patch scores, but they are usually used as filters rather than as signals that reshape the self-supervised objective (Pereira et al., 2019; Liu et al., 2025). Binary filtering discards continuous reliability structure, while static fusion weights ignore cross-modal asymmetry. A useful policy must therefore decide which regions to mask, at what temporal scale, and which modality should provide context, with all three decisions conditioned on the same evidence.

Fixed policies cannot adapt granularity to corruption type while also redistributing fusion weight toward the cleaner modality. To our knowledge, prior work has not conditioned this full pretraining policy on physiological corruption evidence.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

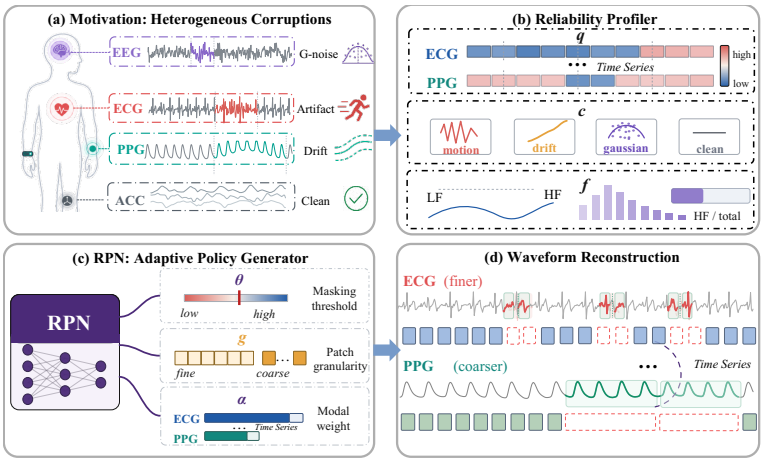

Figure 1: RAMP converts quality, corruption-type, and spectral evidence into an adaptive masking-and-fusion policy that emphasizes cleaner modalities under asymmetric degradation.

Existing methods address pieces of this challenge independently. ST-MEM (Na et al., 2024) models spatio-temporal lead relationships but keeps its masking threshold and granularity fixed. TimeMosaic (Ding et al., 2026) adapts granularity to information density rather than signal quality, and frequency-domain masking (Fu & Hu, 2025) or difficulty-guided sampling (Dong et al., 2023) does not jointly model corruption type, scale, and cross-modal reliability. Static multimodal fusion consequently leaves useful complementarity unused when quality is imbalanced.

We propose RAMP, which learns this joint policy from structured corruption evidence. A Reliability Profiler extracts per-patch quality, corruption type, and spectral degradation; a Reliability Policy Network maps them to coupled masking, granularity, and fusion actions through Gumbel relaxations (Jang et al., 2017). The corruption-type stream is calibrated from a small labeled set (Appendix C.4), while the reconstruction objective remains self-supervised. The policy is used only during pretraining, so downstream transfer uses the same encoder without test-time adaptation. We evaluate on MIMIC-III and VitalDB, with the broader wearable setting left for future validation. Our contributions are fourfold:

- Jointly conditioning masking threshold, temporal granularity, and cross-modal fusion on patch-level corruption evidence yields a single reliability-aware pretraining decision process.
- The Reliability Profiler exposes quality, corruption type, and spectral degradation as complementary streams for policy learning.
- A compact shared Reliability Policy Network couples the three actions through differentiable selection, while RWR preserves informative degraded targets in the reconstruction objective.
- Across ten-seed reconstruction, transfer, ablation, and asymmetric-corruption studies, the joint design improves recovery and representation reuse across six downstream benchmarks.

2 Related Work

SSL for Physiological Time Series. Contrastive methods (Chen et al., 2020; Eldele et al., 2021), masked reconstruction approaches (Zhang et al., 2023), and supervised architectures such as Medformer (Wang et al., 2024) adapt representation learning to physiological time series through domain-specific masking and attention designs. Foundation models (UniTS (Gao et al., 2024), MOMENT (Goswami et al., 2024), Time-MoE (Shi et al., 2025)) generalize across domains; modality-specific variants (Na et al., 2024; Jiang et al., 2024; Wang et al., 2026; Thapa et al., 2024; McKeen et al., 2025; Pillai et al., 2025) scale to large physiological corpora. These works provide effective

encoder architectures and SSL objectives; RAMP addresses a complementary gap: how patches are selected and fused during pretraining when reliability varies within and across modalities.

Adaptive Masking and Multi-Scale Representations. Difficulty-based masking (Dong et al., 2023) samples hard patches by reconstruction error. This criterion is useful for curriculum design, but it can conflate corruption with naturally complex signal morphology: a clinically significant arrhythmia and a motion artifact may both be difficult to reconstruct yet require different treatment. TimeMosaic (Ding et al., 2026) adapts patch granularity to local information density, conditioning on structural complexity rather than per-patch reliability. Frequency-domain masking (Fu & Hu, 2025) and denoising-aware contrastive learning (Zhou et al., 2024) address noise at the objective level, providing complementary strategies for robust representation learning. SleepFM (Thapa et al., 2024) combines multiple physiological modalities, whereas T3Time (Chowdhury et al., 2026) aligns temporal, spectral, and prompt branches for forecasting. RAMP instead focuses on reliability imbalance within the same multimodal physiological recording window. Relative to these adaptive and multi-scale approaches, RAMP conditions the masking policy on observable physiological corruption evidence and coordinates threshold, granularity, and cross-modal weight within a single policy.

Reliability-Conditioned Learning. Signal quality indicators serve as preprocessing filters (Pereira et al., 2019; Liu et al., 2025) or post-hoc data selection tools—not as dynamic conditioning signals that reshape the SSL objective during pretraining. Filtering discards corrupted data before learning; RAMP instead retains corrupted patches as reconstruction targets and adapts the masking policy to draw context from available reliable signal. RelCon (Xu et al., 2025) uses learned semantic similarity to construct contrastive pairs for wearable accelerometry but is limited to segment-level, single-modality settings without per-patch or cross-modal adaptation. Curriculum learning (Bengio et al., 2009) schedules training difficulty globally across examples; it cannot address intra-window patch heterogeneity where artifact and clean morphology coexist within the same recording segment. The information bottleneck perspective motivates the compact evidence state (Tishby et al., 1999; Hu et al., 2024); Gumbel relaxations make the discrete policy actions trainable end-to-end (Jang et al., 2017).

3 METHODOLOGY

3.1 PROBLEM FORMULATION

Let $\mathcal{S} = \{\text{ECG}, \text{PPG}\}$. For each modality $m \in \mathcal{S}$ we observe a segment $\mathbf{x}^m \in \mathbb{R}^T$ partitioned into N non-overlapping patches. Signal reliability varies with corruption type $y_c^m \in \mathcal{C} = \{\text{Gaussian}, \text{Motion}, \text{Drift}, \text{Clean}\}$ (distinct from the corruption-type embedding $\mathbf{c}^m \in \mathbb{R}^K$ defined in Section 3.2). A masking policy π assigns each patch to either the visible context set $\mathcal{V}$ or the masked target set $\mathcal{M}_t$. The pretraining objective minimizes reconstruction error over masked patches:

$$\mathcal{L}(\pi) = \mathbb{E}\left[\sum_{m\in\mathcal{S}} \frac{1}{|\mathcal{M}_t^m|} \sum_{i\in\mathcal{M}_t^m} \|\hat{\mathbf{x}}_i^m - \mathbf{x}_i^m\|_2^2\right], \tag{1}$$

where $\hat{\mathbf{x}}_i^m$ is the decoder prediction conditioned on $\mathcal{V}$ and available cross-modal context (the RWR loss in Eq. 4 is a reliability-weighted generalization; setting $w_i^m \equiv 1$ recovers Eq. 1). Corruption is synthetically injected during pretraining (Appendix B); the reconstruction target $\mathbf{x}_i^m$ is always the *original signal before corruption injection*, so RMSE measures waveform recovery relative to the clean reference.

3.2 MULTI-EVIDENCE RELIABILITY PROFILING

Rather than relying on a scalar quality score, the **Reliability Profiler** (RP) produces a three-component evidence state per modality (Figure 2a):

$$\boldsymbol{\varepsilon}^m = \left[q^m,\ \mathbf{c}^m,\ f^m\right], \tag{2}$$

where $q^m \in [0,1]^N$ is the per-patch reliability score, $\mathbf{c}^m \in \mathbb{R}^K$ ($K = 4$) is a corruption-type embedding, and $f^m \in [0,1]$ is a high-frequency degradation ratio. The Fine-Grained Quality Assessment (FGQA) module estimates q^m under weak segment-level supervision, propagating coarse

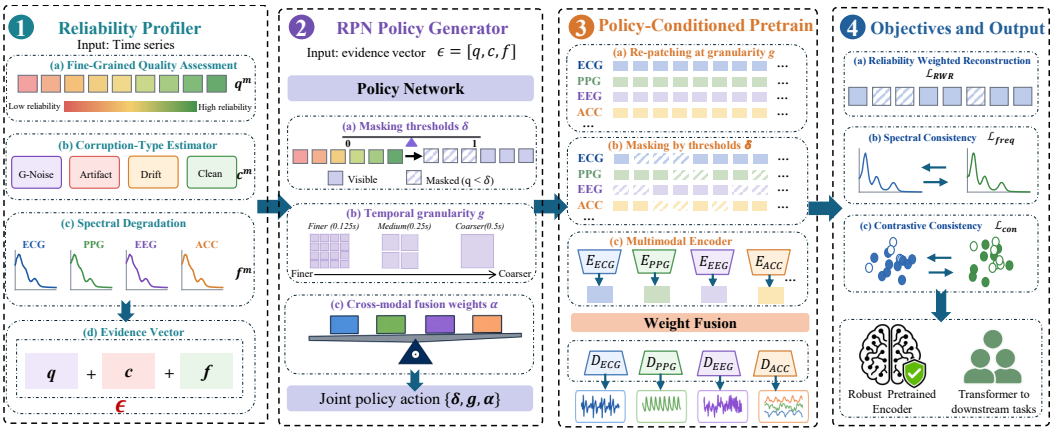

Figure 2: RAMP pretraining extracts per-modality evidence $\boldsymbol{\varepsilon}$, maps it to policy actions δ, g, and $\boldsymbol{\alpha}$, and trains a policy-conditioned reconstruction backbone for downstream transfer.

quality labels to patch-level scores via attention-weighted pooling (Appendix A.1). A corruption-type head predicts Gaussian noise, motion artifact, baseline drift, or clean state; the parameter-free ratio $f^m = P_{>15\text{Hz}}/P_{\text{total}}$ captures broadband degradation, with cutoff sensitivity reported in Appendix A. To validate that FGQA scores reflect signal quality rather than policy preferences, we compute their rank correlation with independent expert ratings: Spearman $\rho = 0.81$ for ECG and 0.78 for PPG (Appendix A). For RPN input, patch scores are summarized by $\bar{q}^m = \text{mean}(\{q_i^m\})$. The segment-level evidence state is $\boldsymbol{\varepsilon} = [\bar{q}^{\text{ECG}}, \mathbf{c}^{\text{ECG}}, f^{\text{ECG}}, \bar{q}^{\text{PPG}}, \mathbf{c}^{\text{PPG}}, f^{\text{PPG}}] \in \mathbb{R}^{2(K+2)}$, with per-patch scores q^m still used for patch-level mask generation.

3.3 Reliability Policy Network

The RPN $\pi_\phi : \boldsymbol{\varepsilon} \mapsto \{\delta, g, \boldsymbol{\alpha}\}$ is a compact MLP (≈12K parameters) that jointly outputs three masking actions: masking threshold $\delta \in [0.4, 0.9]$, which controls the fraction of patches masked via $q^m < \delta$; patch granularity $g \in \{0.125, 0.25, 0.5, 1.0\}$ s, the temporal resolution of mask tokens; and per-modality reliability weights $\boldsymbol{\alpha} = [\alpha^{\text{ECG}}, \alpha^{\text{PPG}}] \in [0, 1]^2$ for reliability-weighted cross-modal fusion. For 5-second, 125-Hz segments, patch length is $P(g) = \lfloor gF_s \rfloor$ samples and token count is $N_g = \lfloor T_s/g \rfloor \in \{5, 10, 20, 40\}$. The resulting mask ratio $r = \mathbb{E}[\mathbb{1}[q_i^m < \delta]]$ is not fixed by design: the RPN jointly learns δ and g so that reconstruction difficulty remains productive; observed mask ratios across corruption types range from 32% to 71% (Appendix C).

Evidence-action alignment. The intended alignment is intuitive: low $\bar{q}^m$ raises δ, motion favors fine g, drift favors coarse g, and high f^m shifts $\boldsymbol{\alpha}$ toward the cleaner modality. It is learned rather than hard-coded, as tested by the evidence ablations in Section 4.

Differentiable policy relaxations. The masking gate applies a Gumbel-sigmoid relaxation of $\mathbb{1}[q_i^m < \delta]$:

$$\hat{a}_i = \sigma\left(\frac{\delta - q_i^m + u - v}{\tau}\right), \quad u, v \sim \text{Gumbel}(0, 1), \tag{3}$$

so $\hat{a}_i$ is high for low-quality patches. The forward pass uses the hard partition and straight-through gradients; granularity uses a straight-through Gumbel-softmax over four values, and a two-way softmax keeps $\alpha^{\text{ECG}} + \alpha^{\text{PPG}} = 1$. Temperature annealing and bias checks are reported in Appendix A.

3.4 Policy-Conditioned Pretraining Backbone

We instantiate Figure 2 with masked reconstruction. At step t, signals are re-patched at g_t, patches with $q^m < \delta_t$ receive mask tokens, and visible patches pass through four convolutional stages, Efficient Channel Attention (Wang et al., 2020), and a 6-layer Transformer (details in Appendix B). Cross-modal representations are fused via $Z_{\text{fused}} = \alpha_t^{\text{ECG}} Z^{\text{ECG}} + \alpha_t^{\text{PPG}} Z^{\text{PPG}}$; Z_{fused} supplies both lightweight modality-specific decoders.

Each pretraining step follows one evidence-to-action loop. The profiler first computes the quality, corruption-type, and spectral state from the corrupted views; the RPN then converts this state into δ_t, g_t, and $\boldsymbol{\alpha}_t$. The hard forward pass re-patches the signals, assigns mask tokens, and fuses the visible representations before either decoder reconstructs the corresponding original signal before corruption injection. The reconstruction and auxiliary losses update the profiler, encoder, decoders, and policy jointly, while the straight-through relaxations provide surrogate gradients for the discrete mask and granularity choices. After pretraining, the profiler and policy are discarded and downstream transfer uses the encoder alone, so adaptation changes the learned representation without introducing test-time policy inference.

Reliability-Weighted Reconstruction (RWR). The RWR loss assigns larger reconstruction weights to masked low-reliability patches, concentrating gradient on recovering degraded signal from clean context:

$$\mathcal{L}_{\text{RWR}} = \frac{\sum_{m\in\mathcal{S}} \sum_{i\in\mathcal{M}_t^m} w_i^m \, \|\hat{\mathbf{x}}_i^m - \mathbf{x}_i^m\|_2^2}{\sum_{m\in\mathcal{S}} \sum_{i\in\mathcal{M}_t^m} w_i^m + \epsilon}, \tag{4}$$

where $w_i^m = 1 + \lambda_{\text{rwr}}(1 - q_i^m)$, $\lambda_{\text{rwr}} = 0.5$, and $\epsilon = 10^{-6}$. Up-weighting achieves lower RMSE than down-weighting, Huber loss, or patch exclusion (Appendix C); corrupted patches retain learnable signal structure. The spectral consistency term penalizes frequency-domain deviation: $\mathcal{L}_{\text{freq}} = \| \, |\hat{X}_m| - |X_m| \, \|_1$, where $|\cdot|$ denotes the short-time Fourier magnitude. The contrastive term enforces segment-level consistency via InfoNCE (Oord et al., 2018):

$$\mathcal{L}_{\text{con}} = -\mathbb{E}\left[\log \frac{e^{\text{sim}(z,z^+)/\tau_{\text{con}}}}{e^{\text{sim}(z,z^+)/\tau_{\text{con}}} + \sum_{j\neq+} e^{\text{sim}(z,z_j^-)/\tau_{\text{con}}}}\right], \tag{5}$$

where z^+ is an augmented view of the same segment, z_j^- are in-batch negatives from different subjects, and $\tau_{\text{con}} = 0.07$ is the contrastive temperature (distinct from the Gumbel temperature τ in Eq. 3). The complete loss is:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{RWR}} + \lambda_f \mathcal{L}_{\text{freq}} + \lambda_c \mathcal{L}_{\text{con}} + \lambda_\varepsilon \|\phi\|_2^2. \tag{6}$$

$\mathcal{L}_{\text{freq}}$ and $\mathcal{L}_{\text{con}}$ serve as standard auxiliary SSL regularizers.

3.5 Design Rationale: Compact Evidence State

The three-stream state $\boldsymbol{\varepsilon} = [q, \mathbf{c}, f]$ is a compact reliability bottleneck: it retains information useful for masking, temporal scale, and fusion without exposing the policy to the full raw signal. The data-processing inequality bounds the policy's predictive information by that in $\boldsymbol{\varepsilon}$; accordingly, the ℓ_2 capacity term in Eq. 6 is used as a practical surrogate rather than as an estimated mutual-information penalty. The formal objective and capacity controls are given in Appendix A.

Training begins with a warm-up phase in which the RPN is clamped to fixed-policy defaults to prevent premature policy collapse, followed by Gumbel temperature annealing (Algorithm 1, Appendix A).

4 Experiments

4.1 Experimental Protocol

The evaluation asks whether reliability-conditioned actions improve reconstruction, whether each evidence stream and action is necessary, and whether the learned encoder transfers beyond ECG/PPG pretraining. We compare fixed and adaptive masking policies with SimCLR, TS-TCC, MAE, MaeFE, SimMTM, ST-MEM, Medformer, and UniTS under matched preprocessing, subject-level splits, optimizer family, and training budgets. Reconstruction uses RMSE, PRD, and SNR; SBP-MAE evaluates linear-probe regression, while six downstream benchmarks report ACC, macro-F1, ROC-AUC, and PR-AUC. Unless noted, results are ten-seed mean$\pm$SD from 100 pretraining epochs and 100 downstream epochs; model and hyperparameter details are given in Appendix B.2. This protocol separates policy quality, representation quality, and transfer utility without changing the encoder family between fixed-policy controls and RAMP.

Table 1: Ten-seed mean ± standard-deviation reconstruction results on MIMIC-III and VitalDB.

	ECG Reconstruction			PPG Reconstruction			SBP Proxy
Method	**RMSE↓**	**PRD (%)↓**	**SNR (dB)↑**	**RMSE↓**	**PRD (%)↓**	**SNR (dB)↑**	**SBP-MAE↓**
Random Masking	0.76 ± 0.10	17.4 ± 1.82	15.1 ± 0.60	0.82 ± 0.11	18.7 ± 1.94	14.3 ± 0.62	4.58 ± 0.19
Contiguous Block Masking	0.88 ± 0.19	19.7 ± 2.55	13.8 ± 0.46	0.95 ± 0.21	21.2 ± 2.71	13.1 ± 0.48	4.83 ± 0.22
Hybrid Masking	0.82 ± 0.15	18.4 ± 1.37	14.7 ± 0.53	0.88 ± 0.17	19.8 ± 1.45	14.0 ± 0.55	4.67 ± 0.20
DM-Rand Masking	0.65 ± 0.08	15.2 ± 1.08	16.4 ± 0.48	0.70 ± 0.09	16.4 ± 1.15	15.6 ± 0.50	4.27 ± 0.17
Fixed Precision ($\delta = 0.7$, $g = 0.5$ s)	0.59 ± 0.07	13.8 ± 0.72	17.3 ± 0.43	0.63 ± 0.08	14.9 ± 0.77	16.5 ± 0.45	4.09 ± 0.17
Fixed Precision + RWR	0.56 ± 0.07	13.1 ± 0.69	17.8 ± 0.41	0.60 ± 0.07	14.2 ± 0.74	16.8 ± 0.43	3.94 ± 0.16
Learnable-δ (MLP, q only)	0.57 ± 0.07	13.4 ± 0.77	17.7 ± 0.42	0.61 ± 0.08	14.4 ± 0.82	16.9 ± 0.44	3.97 ± 0.16
RPN (q only)	0.55 ± 0.07	12.8 ± 0.70	18.1 ± 0.40	0.59 ± 0.07	13.8 ± 0.75	17.3 ± 0.42	3.89 ± 0.15
RPN ($q + f$)	0.53 ± 0.06	12.5 ± 0.68	18.4 ± 0.40	0.57 ± 0.07	13.5 ± 0.73	17.6 ± 0.42	3.81 ± 0.15
RAMP ($q+c+f$, full)	0.50 ± 0.06	11.9 ± 0.65	18.9 ± 0.39	0.54 ± 0.06	12.8 ± 0.70	18.2 ± 0.41	3.71 ± 0.14
RAMP+RWR	$\mathbf{0.47 \pm 0.06}$	$\mathbf{11.2 \pm 0.61}$	$\mathbf{19.5 \pm 0.38}$	$\mathbf{0.51 \pm 0.06}$	$\mathbf{12.1 \pm 0.65}$	$\mathbf{18.8 \pm 0.40}$	$\mathbf{3.58 \pm 0.13}$
Cross-dataset replication: VitalDB (intraoperative; key baselines only; full results in Appendix C)							
Random Masking	0.72 ± 0.09	16.5 ± 1.76	15.7 ± 0.58	0.77 ± 0.10	17.6 ± 1.88	15.1 ± 0.60	4.34 ± 0.18
Fixed Precision + RWR	0.54 ± 0.07	12.6 ± 0.70	17.9 ± 0.41	0.58 ± 0.07	13.2 ± 0.74	17.2 ± 0.42	3.76 ± 0.15
RAMP+RWR	$\mathbf{0.45 \pm 0.06}$	$\mathbf{10.5 \pm 0.58}$	$\mathbf{19.6 \pm 0.38}$	$\mathbf{0.48 \pm 0.06}$	$\mathbf{11.1 \pm 0.61}$	$\mathbf{19.0 \pm 0.39}$	$\mathbf{3.38 \pm 0.13}$

4.2 Datasets and Split Discipline

We pretrain on MIMIC-III (Johnson et al., 2016) (1,000 ICU subjects, 21,473 ECG/PPG segments; 80/10/10 subject split). VitalDB (Lee et al., 2022) (6,388 subjects, 45,213 segments) supports the cross-dataset replication in Table 1: its 90% training split augments pretraining and its held-out 10% provides the evaluation rows, all splits subject-disjoint from MIMIC-III. Both datasets are sampled at 125 Hz in 5-second windows. Downstream evaluation spans six benchmarks across ECG, PPG, EEG, and accelerometer modalities: Epilepsy (Andrzejak et al., 2001), HAR (Anguita et al., 2013), ECG-SQA (Liu et al., 2025), Sleep-EDF (Kemp et al., 2000), PPG-SQA (Liu et al., 2024), and PTB-XL (Wagner et al., 2020) (12-lead diagnostic classification, 5 superclasses; all 12 leads used as input channels). For PPG-SQA, all MIMIC-III subjects and records are disjoint from the pretraining cohort. Full dataset statistics, preprocessing, and the corruption-injection protocol are in Appendix B. An IRB-approved in-house wearable cohort (IHC, 20 subjects) serves as an external feasibility cohort for wearable transfer (Appendix B.1), while large-scale quantitative analyses are primarily based on MIMIC-III and VitalDB.

The split discipline is subject-disjoint at every stage. The MIMIC-III and VitalDB pretraining pools are separated from downstream subjects, and the PPG-SQA evaluation explicitly excludes all pretraining records. This prevents adjacent windows from the same individual from making reconstruction or transfer appear stronger than it is.

4.3 Pretraining Reconstruction

The reconstruction experiment is the primary test of the proposed policy. It holds the encoder family and training budget fixed, then varies only the masking/fusion strategy and the reliability-weighted loss. The top block of Table 1 measures the complete MIMIC-III comparison; the VitalDB block tests whether the ranking survives a change in acquisition setting.

Table 1 shows that RAMP+RWR achieves ECG RMSE 0.47—a 38% reduction over Random Masking and 16% over the strongest fixed-policy baseline (Fixed Precision+RWR)—with comparable absolute gains on PPG. The advantage is consistent across RMSE, PRD, and SNR, and it persists in the VitalDB replication despite the different acquisition setting. This separation matters because it tests both signal recovery and the quality of the learned representation rather than a single reconstruction metric.

Every quality-agnostic strategy trails its quality-aware counterpart; even DM-Rand, which oversamples hard patches, does not eliminate the performance gap. Contiguous Block masking underperforms Random Masking (ECG RMSE 0.88 vs. 0.76): excessively long unmasked spans deprive the decoder of the contextual diversity necessary to reconstruct structured physiological patterns. On VitalDB, RAMP+RWR achieves ECG RMSE 0.45 and PPG RMSE 0.48, corroborating the MIMIC-III ranking with wider absolute margins—consistent with the elevated baseline noise characteristic

Table 2: Ablation results for reconstruction and transfer.

(a) End-to-end contribution

Variant	Design	ECG↓	PPG↓	SBP↓	Avg. ACC↑	Δ PPG
Fixed Prec.	fixed δ, g, α	0.59	0.63	4.09	93.2	+0.12
Fixed Prec. + RWR	fixed policy + RWR	0.56	0.60	3.94	93.4	+0.09
RAMP (no RWR)	joint $q + c + f$ policy	0.50	0.54	3.71	93.6	+0.03
RAMP+RWR	**joint policy + RWR**	**0.47**	**0.51**	**3.58**	**93.9**	**—**

(b) Evidence and action controls (no RWR)

Variant	Intervention	PPG		Δ vs. full	Question tested
		mean	change	value	
Full RPN (q+c+f)	reference	**0.54**	—	—	—
Shuffled **c**	remove c alignment	0.57	+0.03	+0.03	granularity
Random f	remove f alignment	0.56	+0.02	+0.02	modality
Shuffled q	permute q	0.61	+0.07	+0.07	quality
q only	remove c, f	0.59	+0.05	+0.05	complement
All random	capacity control	0.63	+0.09	+0.09	capacity
Adaptive δ only	fix g, α	0.61	+0.07	+0.07	threshold
Adaptive g only	fix δ, α	0.60	+0.06	+0.06	scale
Adaptive α only	fix δ, g	0.62	+0.08	+0.08	fusion
Adaptive $\delta + g$	fix α	0.57	+0.03	+0.03	paired actions
Separate heads	decouple actions	0.59	+0.05	+0.05	coupling

Table 3: Reconstruction RMSE by corruption type and modality.

		Fixed masking policies						Adaptive policy			
Corruption	**Mod.**	Rand.	Cont.	Hybrid	DM-Rand	Fixed	Learn.-δ	RPN(q)	RPN(q+f)	RAMP	**+RWR**
Clean	E	0.41	0.51	0.46	0.38	0.34	0.33	0.31	0.30	0.28	**0.26**
	P	0.44	0.55	0.49	0.41	0.36	0.35	0.33	0.32	0.30	**0.28**
Gaussian	E	0.68	0.82	0.75	0.61	0.52	0.50	0.48	0.46	0.43	**0.40**
	P	0.73	0.88	0.80	0.65	0.55	0.53	0.51	0.49	0.46	**0.43**
Motion	E	0.91	1.09	0.98	0.79	0.67	0.65	0.61	0.58	0.54	**0.50**
	P	0.97	1.16	1.05	0.84	0.71	0.69	0.65	0.62	0.57	**0.53**
Drift	E	0.57	0.68	0.63	0.51	0.43	0.42	0.40	0.38	0.36	**0.33**
	P	0.61	0.73	0.67	0.54	0.46	0.45	0.43	0.41	0.39	**0.36**
Mixed	E	0.94	1.11	1.02	0.82	0.70	0.68	0.64	0.62	0.58	**0.54**
	P	0.99	1.18	1.08	0.87	0.74	0.72	0.68	0.65	0.61	**0.57**

of intraoperative recordings. Figure 3 visualizes representative RPN decisions and their effect on recovered waveforms across corruption types. Table 3 gives the detailed corruption-by-modality comparison discussed in Section 4.6.

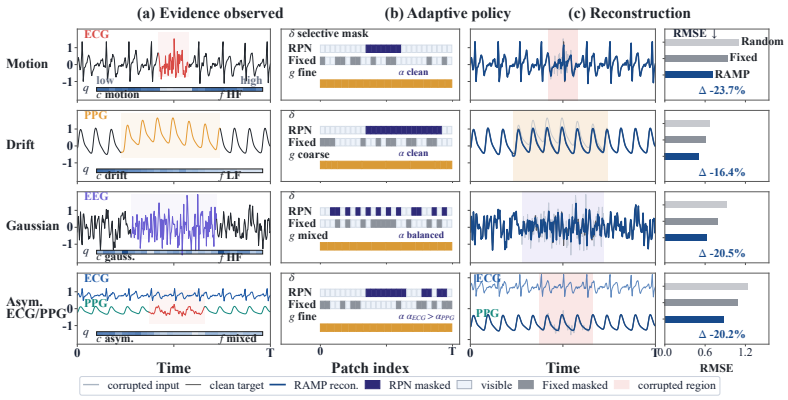

Figure 3: Reliability evidence drives adaptive masks and recovered waveforms.

4.4 Evidence and Action Ablations

Table 2 separates the end-to-end contribution from the evidence-to-action controls. The end-to-end block shows complementary gains: RWR improves the fixed-policy baseline from 0.63 to 0.60 PPG RMSE, while the joint policy without RWR reaches 0.54; combining both reaches 0.51 and also

reduces SBP-MAE from 4.09 to 3.58. The corresponding downstream average rises from 93.2% for Fixed Precision to 93.6% for the policy alone and 93.9% for the full model, indicating that the reconstruction and transfer effects are not interchangeable.

The control block identifies which evidence is used. Shuffling $\mathbf{c}$ or randomizing f produces smaller but consistent degradations, whereas shuffling q or replacing all evidence produces larger losses. On the action side, each single-action policy remains below the joint policy, and separate action heads increase RMSE by 0.05. These controls preserve input dimensionality while breaking either sample-level evidence correspondence or cross-action coupling, supporting the interpretation that the gain arises from coordinated evidence use rather than additional parameters. Appendix C reports the weighting, stage-depth, data-efficiency, and objective-specific controls. Figure 4 summarizes the evidence-stream and action-level controls discussed in this subsection.

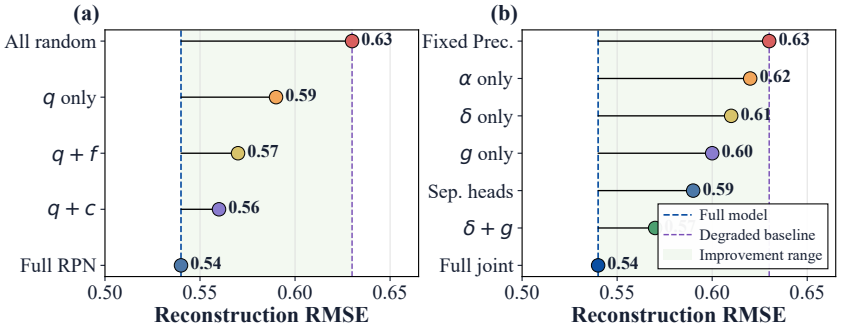

Figure 4: Evidence-stream and action-level ablations show the PPG RMSE gap from the full RPN policy.

4.5 Downstream Transfer and Statistical Validation

The learned encoder is frozen or linearly probed for downstream tasks, so the policy is not re-optimized after pretraining. Table 4 reports ACC, macro-F1, and ROC-AUC for six benchmarks spanning EEG, accelerometry, ECG, PPG, and twelve-lead diagnostic classification; the full mean±SD decomposition and PR-AUC results are in Appendix C. This protocol tests whether reliability-aware reconstruction produces reusable features rather than merely lowering training-time signal error.

The aggregate table shows a consistent transfer pattern. Relative to Fixed Precision, RAMP+RWR improves accuracy on Epilepsy, HAR, ECG-SQA, Sleep-EDF, and PTB-XL, while ST-MEM remains strongest on PPG-SQA. The gains are not restricted to accuracy: HAR shows the largest ROC-AUC margin, and the appendix reports the corresponding PR-AUC changes and ten-seed standard deviations. The policy and loss contributions are also separable: the policy raises mean ACC from 93.2% to 93.6% before RWR, and RWR contributes the remaining 0.3 percentage points.

Paired Wilcoxon tests on per-seed accuracy give $p = 0.033$ for HAR, $p = 0.021$ for ECG-SQA and PTB-XL, $p = 0.084$ for Epilepsy, $p = 0.063$ for PPG-SQA, and $p = 0.107$ for Sleep-EDF. The pattern is therefore strongest on tasks whose input distributions contain heterogeneous quality or diagnostic corruption, while the Sleep-EDF result remains a small, non-significant gain. This distinction is important: the transfer evidence supports broad feature reuse, but does not imply uniform improvement on every target dataset.

4.6 Corruption-Specific Reconstruction

When signal quality diverges across modalities, fixed fusion weights become systematically miscalibrated. Table 3 makes this comparison explicit across clean, Gaussian, motion, drift, and mixed corruption for both ECG and PPG. RAMP+RWR obtains the lowest RMSE in all ten conditions and improves on RAMP without RWR by 0.02–0.04 RMSE, with the largest margins for Motion and Mixed corruption. The ordering is consistent across modalities, indicating that the gain is not tied to a single signal type.

Table 4: Downstream classification across six benchmarks reports mean ACC / macro-F1 / ROC-AUC over ten seeds (%).

Method	**Epilepsy**	**HAR**	**ECG-SQA**	**S-EDF**	**PPG-SQA**	**PTB-XL**
	each cell: ACC / macro-F1 / AUC-ROC (%)					
SimCLR	94.7/94.0/93.2	91.8/91.3/91.0	91.3/90.2/90.3	81.9/77.2/78.9	90.6/90.3/90.8	91.8/81.8/90.2
MAE	96.4/96.5/95.3	93.0/92.1/92.0	92.2/91.2/91.0	81.8/79.0/80.0	91.3/93.8/92.0	92.0/87.8/87.4
MaeFE	96.6/96.6/96.3	92.8/92.3/92.5	90.4/90.4/90.3	77.6/76.3/76.9	93.2/92.8/92.0	89.0/88.6/88.3
TS-TCC	95.7/94.9/95.0	92.8/92.3/92.0	88.0/87.9/87.4	80.7/80.2/79.9	94.6/94.4/93.2	91.7/92.4/89.8
ST-MEM	96.7/96.0/94.2	92.5/92.4/92.5	87.8/88.8/88.7	78.0/77.0/78.0	**98.2**/92.3/92.0	91.0/87.7/87.9
SimMTM	94.4/93.5/94.1	93.3/92.1/92.1	90.9/90.6/90.8	79.0/77.0/76.3	94.3/93.9/93.8	88.1/86.4/85.3
Medformer	97.5/97.3/96.3	91.5/91.4/90.6	89.3/89.0/88.3	77.7/78.7/75.5	92.2/93.0/92.0	89.8/89.4/86.6
Foundation model — linear probe only; scale context						
UniTS	94.5/93.6/94.3	91.5/91.0/90.4	88.9/88.5/88.2	77.1/76.0/75.6	94.8/94.5/94.3	90.0/87.4/86.6
Fixed Prec.	97.8/**97.8**/97.8	94.3/93.5/90.6	93.5/93.0/92.3	84.3/80.3/82.3	96.5/**95.9**/**96.0**	92.5/91.7/92.3
RAMP (no RWR)	98.0/97.6/97.9	94.8/94.2/93.1	94.2/93.9/93.0	84.4/84.5/83.8	96.7/95.7/95.4	93.2/92.6/92.7
RAMP+RWR	**98.1**/97.5/**98.0**	**95.2**/**94.7**/**95.0**	**94.8**/**94.6**/**93.5**	**84.5**/**87.6**/**84.9**	96.9/95.5/94.9	**93.8**/**93.3**/**93.0**

The corruption-specific pattern also clarifies the mechanism. Motion and Mixed inputs benefit most from combining spectral and corruption-type evidence because their reliable context changes within a window; drift is less abrupt and is better served by a coarser action. RWR then keeps degraded but informative targets in the objective without allowing them to dominate the update. Thus the result is not explained by simply masking more aggressively: the policy chooses where and at what scale to reconstruct, while RWR controls how those targets contribute to learning. Policy-output dynamics are reported in Appendix C.

5 Discussion

The central result is a coordination effect. Reliability evidence is useful not only for ranking patches, but also for deciding the temporal scale of the reconstruction target and the modality that supplies context. The policy-only and RWR-only controls show that these mechanisms address different failure modes: policy selection changes the learning problem, while RWR changes how gradients are allocated across the selected targets. The asymmetric analysis further shows that the learned behavior follows the reliability structure rather than a single global masking preference. This complementarity is consistent across the MIMIC-III and VitalDB evaluations under the matched-backbone protocol, including asymmetric corruption; encoder-size controls further support policy conditioning as the source of the gain.

The evidence remains scoped to controlled corruption, ICU and intraoperative data, and a pretraining-time policy. The appendix establishes robustness to unseen synthetic corruption, calibration noise, data scale, and encoder size, but naturally corrupted wearables, missing modalities, and online recalibration remain open evaluation settings. These boundaries motivate the next step: calibrating the evidence state on sensor-specific natural artifacts while preserving the subject-disjoint transfer protocol.

6 Conclusion

RAMP reframes masked pretraining as a reliability-conditioned decision process rather than a fixed corruption schedule. Its profiler represents patch quality, corruption type, and spectral degradation, while a shared policy coordinates masking thresholds, temporal granularity, and cross-modal fusion. Reliability-Weighted Reconstruction complements these actions by retaining degraded targets that still carry physiological structure. Across subject- disjoint MIMIC-III and VitalDB evaluations, this coordination improves reconstruction under matched and asymmetric corruption and transfers across ECG, PPG, EEG, and accelerometry tasks. The evidence remains bounded to controlled corruption and pretraining-time adaptation; extending the evidence state to natural wearable artifacts, missing modalities, and online calibration is the next step toward deployment-oriented evaluation.

AI USE STATEMENT

Generative AI tools assisted with code development; the authors verified the analyses and final manuscript.

ETHICS STATEMENT

The IHC study received approval from an institutional ethics committee; the institution name and approval identifier are withheld during double-blind review. All IHC data were de-identified before analysis. The remaining evaluations use de-identified public datasets.

REPRODUCIBILITY STATEMENT

The supplement specifies the information boundary, dataset mapping, participant-level splits, feature construction, model selection, baselines, robustness protocols, statistical analyses, and computational resources. The anonymous implementation is available at `https://anonymous.4open.science/r/RAMP-Physio/`.

REFERENCES

Ralph G Andrzejak, Klaus Lehnertz, Florian Mormann, Christoph Rieke, Peter David, and Christian E Elger. Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Physical Review E*, 64(6):061907, 2001. doi: 10.1103/PhysRevE.64.061907.

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pp. 437–442, 2013.

Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15619–15629, 2023. doi: 10.1109/CVPR52729.2023.01499.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009. doi: 10.1145/1553374.1553380.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020.

Abdul Monaf Chowdhury, Rabeya Akter, and Safaeid Hossain Arib. T3Time: Tri-modal time series forecasting via adaptive multi-head alignment and residual fusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 20597–20605, 2026. doi: 10.1609/aaai.v40i25.39196.

Kuiye Ding, Fanda Fan, Chunyi Hou, Zheya Wang, Lei Wang, Zhengxin Yang, and Jianfeng Zhan. TimeMosaic: Temporal heterogeneity guided time series forecasting via adaptive granularity patch and segment-wise decoding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 20790–20798, 2026. doi: 10.1609/aaai.v40i25.39218.

Jiaxiang Dong, Haixu Wu, Haoran Zhang, Li Zhang, Jianmin Wang, and Mingsheng Long. SimMTM: A simple pre-training framework for masked time-series modeling. *Advances in Neural Information Processing Systems*, 36:29996–30025, 2023. doi: 10.52202/075280-1306.

Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, Chee Keong Kwoh, Xiaoli Li, and Cuntai Guan. Time-series representation learning via temporal and contextual contrasting. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pp. 2352–2359, 2021. doi: 10.24963/ijcai.2021/324.

En Fu and Yanyan Hu. Frequency-masked embedding inference: a non-contrastive approach for time series representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 16639–16647, 2025. doi: 10.1609/aaai.v39i16.33828.

Sanjiv S Gambhir, T Jessie Ge, Ophir Vermesh, Ryan Spitler, and Garry E Gold. Continuous health monitoring: An opportunity for precision health. *Science translational medicine*, 13(597): eabe5383, 2021. doi: 10.1126/scitranslmed.abe5383.

Shanghua Gao, Teddy Koker, Owen Queen, Thomas Hartvigsen, Theodoros Tsiligkaridis, and Marinka Zitnik. UniTS: A unified multi-task time series model. *Advances in Neural Information Processing Systems*, 37, 2024.

Mononito Goswami, Konrad Szafer, Arjun Choudhry, Yifu Cai, Shuo Li, and Artur Dubrawski. MOMENT: A family of open time-series foundation models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 16115–16152, 2024.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022. doi: 10.1109/CVPR52688.2022.01553.

Shizhe Hu, Zhengzheng Lou, Xiaoqiang Yan, and Yangdong Ye. A survey on information bottleneck. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5325–5344, 2024. doi: 10.1109/TPAMI.2024.3366349.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*, 2017. Published at ICLR 2017; arXiv preprint posted in 2016.

Wei-Bang Jiang, Liming Zhao, and Bao-Liang Lu. Large brain model for learning generic representations with tremendous EEG data in BCI. In *International Conference on Learning Representations*, 2024.

Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):160035, 2016. doi: 10.1038/sdata.2016.35.

Bob Kemp, Aeilko H Zwinderman, Bert Tuk, Hilbert AC Kamphuisen, and Josefien JL Oberye. Analysis of a sleep-dependent neuronal feedback loop: the slow-wave microcontinuity of the eeg. *IEEE Transactions on Biomedical Engineering*, 47(9):1185–1194, 2000. doi: 10.1109/10.867928.

Rayan Krishnan, Pranav Rajpurkar, and Eric J Topol. Self-supervised learning in medicine and healthcare. *Nature Biomedical Engineering*, 6(12):1346–1352, 2022. doi: 10.1038/s41551-022-00914-1.

Hyung-Chul Lee, Yoonsang Park, Soo Bin Yoon, Seong Mi Yang, Dongnyeok Park, and Chul-Woo Jung. Vitaldb, a high-fidelity multi-parameter vital signs database in surgical patients. *Scientific Data*, 9(1):279, 2022. doi: 10.1038/s41597-022-01411-5.

Jian Liu, Shuaicong Hu, Ya'nan Wang, Qihan Hu, Daomiao Wang, and Cuiwei Yang. A lightweight hybrid model using multiscale markov transition field for real-time quality assessment of photoplethysmography signals. *IEEE Journal of Biomedical and Health Informatics*, 28(2):1078–1088, 2024. doi: 10.1109/JBHI.2023.3331975.

Jian Liu, Shuaicong Hu, Yanan Wang, Qihan Hu, Daomiao Wang, Wei Xiang, Xujian Feng, and Cuiwei Yang. Leaf-net: A real-time fine-grained quality assessment system for physiological signals using lightweight evolutionary attention fusion. *Expert Systems with Applications*, 274: 126995, 2025. doi: 10.1016/j.eswa.2025.126995.

Kaden McKeen, Sameer Masood, Augustin Toma, Barry Rubin, and Bo Wang. ECG-FM: An open electrocardiogram foundation model. *Jamia Open*, 8(5):ooaf122, 2025. doi: 10.1093/jamiaopen/ooaf122.

Yeongyeon Na, Minje Park, Yunwon Tae, and Sunghoon Joo. Guiding masked representation learning to capture spatio-temporal relationship of electrocardiogram. In *International Conference on Learning Representations*, 2024.

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

Junyung Park, Hyeon Seok Seok, Sang-Su Kim, and Hangsik Shin. Photoplethysmogram analysis and applications: an integrative review. *Frontiers in physiology*, 12:808451, 2022. doi: 10.3389/fphys.2021.808451.

Tania Pereira, Cheng Ding, Kais Gadhoumi, Nate Tran, Rene A Colorado, Karl Meisel, and Xiao Hu. Deep learning approaches for plethysmography signal quality assessment in the presence of atrial fibrillation. *Physiological measurement*, 40(12):125002, 2019. doi: 10.1088/1361-6579/ab5b84.

Arvind Pillai, Dimitris Spathis, Fahim Kawsar, and Mohammad Malekzadeh. PaPaGei: Open foundation models for optical physiological signals. In *International Conference on Learning Representations*, 2025.

Francesco Scardulla, Gloria Cosoli, Susanna Spinsante, Angelica Poli, Grazia Iadarola, Riccardo Pernice, Alessandro Busacca, Salvatore Pasta, Lorenzo Scalise, and Leonardo D'Acquisto. Photoplethysmograhic sensors, potential and limitations: Is it time for regulation? a comprehensive review. *Measurement*, 218:113150, 2023. doi: 10.1016/j.measurement.2023.113150.

Xiaoming Shi, Shiyu Wang, Yuqi Nie, Dianqi Li, Zhou Ye, Qingsong Wen, and Ming Jin. Time-MoE: Billion-scale time series foundation models with mixture of experts. In *International Conference on Learning Representations*, 2025.

Rahul Thapa, Bryan He, Magnus Ruud Kjær, Hyatt Moore, Gauri Ganjoo, Emmanuel Mignot, and James Zou. Sleepfm: multi-modal representation learning for sleep across brain activity, ecg and respiratory signals. *arXiv preprint arXiv:2405.17766*, 2024.

Hugo Thimonier, José Lucas De Melo Costa, Fabrice Popineau, Arpad Rimmel, and Bich-Liên Doan. T-JEPA: Augmentation-free self-supervised learning for tabular data. In *International Conference on Learning Representations*, 2025.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, pp. 368–377, 1999. Conference publication in 1999; arXiv version posted in 2000.

Gaël Vila, Christelle Godin, Sylvie Charbonnier, and Aurélie Campagne. Real-time quality index to control data loss in real-life cardiac monitoring applications. *Sensors*, 21(16):5357, 2021. doi: 10.3390/s21165357.

Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bousseljot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. Ptb-xl, a large publicly available electrocardiography dataset. *Scientific data*, 7(1):154, 2020. doi: 10.1038/s41597-020-0495-6.

Qilong Wang, Banggu Wu, Pengfei Zhu, Peihua Li, Wangmeng Zuo, and Qinghua Hu. Eca-net: Efficient channel attention for deep convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11534–11542, 2020. doi: 10.1109/CVPR42600.2020.01155.

Xin Wang, Burcu Ozek, Aruna Mohan, Amirhossein Ravari, Or Zilbershot, and Fatemeh Afghah. RhythmBERT: A self-supervised language model based on latent representations of ECG waveforms for heart disease detection. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6351–6355. IEEE, 2026. doi: 10.1109/ICASSP55912.2026.11461030.

Yihe Wang, Nan Huang, Taida Li, Yujun Yan, and Xiang Zhang. Medformer: A multi-granularity patching transformer for medical time-series classification. *Advances in Neural Information Processing Systems*, 37:36314–36341, 2024.

Maxwell Xu, Jaya Narain, Gregory Darnell, Haraldur Hallgrimsson, Hyewon Jeong, Darren Forde, Richard Fineman, Karthik Raghuram, James Rehg, and Shirley Ren. RelCon: Relative contrastive learning for a motion foundation model for wearable data. In *International Conference on Learning Representations*, 2025.

Huaicheng Zhang, Wenhan Liu, Jiguang Shi, Sheng Chang, Hao Wang, Jin He, and Qijun Huang. MaeFE: Masked autoencoders family of electrocardiogram for self-supervised pretraining and transfer learning. *IEEE Transactions on Instrumentation and Measurement*, 72:1–15, 2023. doi: 10.1109/TIM.2022.3228267.

Kexin Zhang, Qingsong Wen, Chaoli Zhang, Rongyao Cai, Ming Jin, Yong Liu, James Y Zhang, Yuxuan Liang, Guansong Pang, Dongjin Song, et al. Self-supervised learning for time series analysis: Taxonomy, progress, and prospects. *IEEE transactions on pattern analysis and machine intelligence*, 46(10):6775–6794, 2024. doi: 10.1109/TPAMI.2024.3387317.

Shuang Zhou, Daochen Zha, Xiao Shen, Xiao Huang, Rui Zhang, and Fu-Lai Chung. Denoising-aware contrastive learning for noisy time series. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pp. 5644–5652, 2024. doi: 10.24963/ijcai.2024/624.

RAMP: SUPPLEMENTARY MATERIAL

Anonymous authors
Paper under double-blind review

NOMENCLATURE

A EXTENDED METHODOLOGY

A.1 IMPLEMENTATION DETAILS

The Fine-Grained Quality Assessment (FGQA) module adopts a four-stage hierarchical design alternating Multi-Scale Feature Aggregator (MSFA) blocks and Multi-Head Self-Attention (MSA) blocks. Table A2 summarizes the per-stage configuration. Input is a [B, 1, L] tensor; MSFA blocks apply multi-head convolutions (MHC) at kernel sizes $\{3, 7, 15\}$ with an Efficient Channel Attention (ECA) module, then aggregate across heads.

The corruption-type classifier head is a 2-layer MLP (256→64→4) with GELU activation attached to the Stage IV pooled representation, trained with cross-entropy loss on 500 manually annotated calibration segments per modality. The spectral degradation feature f^m is computed from the Welch PSD estimate (window 0.5 s, 50% overlap) as $f^m = \sum_{k:\, f_k > 15\text{Hz}} S_k \;/\; \sum_k S_k$.

FGQA quality score semantic validity. Spearman rank correlations between per-patch FGQA outputs and independent expert quality ratings confirm that FGQA scores reflect clinically meaningful quality variation independently of the masking policy. The held-out evaluation uses a 200-segment subset from the MIMIC-III test set (disjoint from all training and calibration segments); two biomedical engineers rated each patch on a five-level quality scale normalized to $[0, 1]$ following the protocol of Section C.4, without knowledge of FGQA scores. Inter-rater agreement after binning was weighted Cohen's $\kappa = 0.85$; the resulting Spearman $\rho = 0.81$ (ECG) and $\rho = 0.78$ (PPG) confirm the validity claim.

RPN training. The RPN is initialized with Kaiming uniform weights and trained from epoch 1 alongside the encoder-decoder. During the first 5 warm-up epochs the RPN outputs are clamped to the Fixed-Precision defaults ($\delta = 0.7$, $g = 0.5$ s, $\boldsymbol{\alpha} = [0.5, 0.5]$) to allow stable encoder initialization. The Gumbel temperature τ follows a cosine annealing schedule from 1.0 to 0.1 over the 100 pretraining epochs. The L_2 regularizer weight $\lambda_\varepsilon = 10^{-3}$ was selected from $\{10^{-4}, 10^{-3}, 10^{-2}\}$ via validation-set RMSE (0.53, **0.51**, 0.56 respectively): weak regularization overfits to batch noise; strong regularization collapses the policy toward fixed-threshold behavior.

Pretraining algorithm and fine-tuning. Algorithm 1 details the full pretraining step, combining Reliability Profiling and RPN action selection with the joint loss $\mathcal{L}_{\text{total}}$ (Eq. 6 in the main text).

Algorithm 1 RAMP pretraining step

Require: Batch $\{\mathbf{x}^m\}_{m\in\mathcal{M}}$; FGQA module g_ψ; RPN π_ϕ; Encoder f_ω; Decoders $\{d_\eta^m\}$
1: Extract evidence: $\boldsymbol{\varepsilon} \leftarrow [g_\psi(\mathbf{x}^{\text{ECG}}); g_\psi(\mathbf{x}^{\text{PPG}})]$ {quality, type, spectral}
2: Compute policy: $(\delta, g, \boldsymbol{\alpha}) \leftarrow \pi_\phi(\boldsymbol{\varepsilon})$ via scaled sigmoid / ST-Gumbel-softmax / softmax
3: Generate mask: $\mathcal{M}_t \leftarrow \{i : q_i^m < \delta\}$ with granularity g
4: Encode visible: $Z \leftarrow f_\omega(\mathbf{x}_\mathcal{V})$
5: Reliability-weighted fusion: $Z_{\text{fused}} \leftarrow \alpha^{\text{ECG}} \cdot Z^{\text{ECG}} + \alpha^{\text{PPG}} \cdot Z^{\text{PPG}}$
6: Decode masked: $\hat{\mathbf{x}}^m \leftarrow d_\eta^m(Z_{\text{fused}})$
7: Compute $\mathcal{L}_{\text{total}}$ via Eq. 6 (main paper)
8: Update $\{\psi, \omega, \eta, \phi\}$ by back-propagation

Table A1: Abbreviations used in this paper.

Abbreviation	Full Form
Model components	
RAMP	Reliability-Adaptive Masking Policy
RPN	Reliability Policy Network
RIB	Reliability Information Bottleneck
RWR	Reliability-Weighted Reconstruction
FGQA	Fine-Grained Quality Assessment
MSFA	Multi-Scale Feature Aggregator
MSA	Multi-Head Self-Attention
MHC	Multi-Head Convolution
ECA	Efficient Channel Attention
MLP	Multi-Layer Perceptron
ST	Straight-Through (Gumbel estimator)
Learning paradigms	
SSL	Self-Supervised Learning
MAE*	Masked Autoencoder (also: Mean Absolute Error in metrics)
DM-Rand	Difficulty-Matched Random masking
OOD	Out-of-Distribution
Physiological signals & datasets	
ECG	Electrocardiography
PPG	Photoplethysmography
EEG	Electroencephalography
BP	Blood Pressure
SBP	Systolic Blood Pressure
DBP	Diastolic Blood Pressure
HAR	Human Activity Recognition
SQA	Signal Quality Assessment
PSD	Power Spectral Density
IHC	In-House Cohort
ICU	Intensive Care Unit
Evaluation metrics	
RMSE	Root Mean Square Error
MAE*	Mean Absolute Error (also: Masked Autoencoder above)
ME	Mean Error
SD	Standard Deviation
PRD	Percent Root-mean-square Difference
SNR	Signal-to-Noise Ratio
ACC	Classification Accuracy
F1	Macro-averaged F1 Score
ROC	Receiver Operating Characteristic
PR	Precision-Recall
AUC	Area Under the Curve

Table A2: FGQA stage-wise architecture and parameter count. L_s/C_s: sequence length/channel dimension after downsampling; Stage IV includes the quality and corruption-type output heads. The two modality-specific profilers total 2.506M parameters.

Stage	Blocks	C_s	L_s	Params (K)
I	2×MSFA	32	$L/4$	18.5
II	2×MSFA + 1×MSA	64	$L/8$	94.9
III	2×MSFA + 2×MSA	128	$L/16$	427.4
IV	3×MSA	256	$L/32$	712.3
Total per modality				**1,253.1**
Total for two modalities				**2,506.2**

Downstream fine-tuning. For BP estimation, the encoder is attached to a 2-layer MLP head (512→256→1) and fine-tuned for 50 epochs with Adam (lr $= 10^{-4}$, cosine decay, weight decay 10^{-4}). For classification, a linear layer is appended (linear evaluation) or the full encoder is fine-tuned for 100 epochs; dataset-provided splits are used throughout.

A.2 Design Rationale: Information-Theoretic Perspective

The RIB objective (main paper Eq. 7) provides an interpretive design perspective that balances task relevance against evidence compression:

$$\mathcal{L}_{\mathrm{RIB}}(\phi) = -I\big(\pi_\phi(\varepsilon);\, Z^*\big) + \beta_\varepsilon\, I\big(\pi_\phi(\varepsilon);\, \varepsilon \mid Z^*\big), \tag{A.1}$$

where Z^* is the oracle representation maximally predictive of masked patches. By the data-processing inequality (Tishby et al., 1999), $I(\pi_\phi(\varepsilon); Z^*) \leq I(\varepsilon; Z^*)$, so richer evidence directly raises the policy's representational ceiling. Increasing β_ε favors weaker dependence on the evidence state and therefore approaches the behavior of evidence-independent fixed policies. We do not estimate the conditional mutual information during training; $\lambda_\varepsilon \|\phi\|_2^2$ is instead a practical capacity-control surrogate, with λ_ε selected by validation-set RMSE.

Action-specific evidence requirements. Threshold δ benefits from $\bar{q}^m$ (ranks patch reliability); granularity g benefits from $\mathbf{c}^m$ (corruption type determines temporal scale); fusion weight $\boldsymbol{\alpha}$ benefits from f^m (spectral degradation captures modality-level quality not encoded in per-patch scores). These roles are supported empirically by the action-level ablation (Table A20).

Rationale for joint optimization. The masked set $\mathcal{M}_t$ depends jointly on $(\delta, g, \boldsymbol{\alpha})$, so changing one action alters the context in which the other actions are evaluated. A shared RPN permits gradients from these coupled decisions to update a common evidence representation. The separate-heads ablation (Table A20) raises RMSE from 0.54 to 0.59 ($+0.05$), which is consistent with beneficial shared-parameter coupling; this design perspective does not predict the magnitude of the empirical gain.

A.3 Differentiable Policy: Complete Derivation

A. Patch masking gate $\hat{a}_i$ (Gumbel-sigmoid). The threshold $\delta \in [0.4, 0.9]$ is a continuous RPN output mapped via scaled sigmoid: $\delta = 0.4 + 0.5\,\sigma(h_\delta(\boldsymbol{\varepsilon}))$; no Gumbel relaxation is needed for δ itself. The Gumbel-sigmoid instead relaxes the per-patch binary mask decision $\mathbb{1}[q_i^m < \delta]$. Let $\eta_i = u_i - v_i$ with $u_i, v_i \sim \mathrm{Gumbel}(0,1)$:

$$\hat{a}_i = \sigma\left(\frac{\delta - q_i^m + \eta_i}{\tau}\right), \tag{A.2}$$

where $\sigma(x) = (1+e^{-x})^{-1}$. By construction, $\hat{a}_i \approx 1$ when $q_i^m < \delta$ (masked) and $\hat{a}_i \approx 0$ otherwise. The gradient with respect to η_i satisfies the following bound.

Proposition 1 (Gradient Boundedness)**.** *For any differentiable loss $\mathcal{L}$ and finite $\tau > 0$:*

$$\left|\frac{\partial \mathcal{L}}{\partial \eta_i}\right| \le \frac{1}{4\tau}\left|\frac{\partial \mathcal{L}}{\partial \hat{a}_i}\right|. \tag{A.3}$$

Proof. $\sigma(x)(1-\sigma(x)) \le \frac{1}{4}$ for all $x \in \mathbb{R}$; substituting into the chain-rule expression for $\partial\mathcal{L}/\partial\eta_i$ and dividing by τ gives the bound. □

Remark. The bound peaks at $1/(4\tau_{\min}) = 2.5$ when $\tau = 0.1$ and decreases monotonically toward 0.25 as $\tau \to 1$, preventing gradient explosion while maintaining informative update signal throughout annealing.

B. Patch granularity g (Straight-through Gumbel-softmax). Granularity is selected from $\mathcal{G} = \{g_1, \ldots, g_K\}$ with $K = 4$. Let $\mathbf{u} = (u_1, \ldots, u_K)$ be RPN logits; draw $\boldsymbol{\gamma} \sim \mathrm{Gumbel}(0,1)^K$. The forward pass uses the hard assignment:

$$\hat{k} = \underset{k=1,\ldots,K}{\arg\max}\,(u_k + \gamma_k), \quad g^{\mathrm{sel}} = g_{\hat{k}}. \tag{A.4}$$

The backward pass uses the soft Gumbel-softmax surrogate:

$$\tilde{p}_k = \frac{\exp\big((u_k + \gamma_k)/\tau\big)}{\sum_{j=1}^{K} \exp\big((u_j + \gamma_j)/\tau\big)}, \tag{A.5}$$

and the straight-through (ST) gradient estimate is:

$$\hat{\nabla}_{\mathbf{u}} \mathcal{L} = \big(\nabla_{\tilde{\mathbf{p}}} \mathcal{L}\big)\big|_{\tilde{\mathbf{p}} = \tilde{\mathbf{p}}(\mathbf{u}, \boldsymbol{\gamma})}. \tag{A.6}$$

The ST estimator is biased but commonly used and empirically stable (Jang et al., 2017); finite τ keeps the soft surrogate differentiable, while annealing makes its assignments increasingly concentrated.

C. Cross-modal alignment weight α (softmax, exact). Let $u_\alpha \in \mathbb{R}$ be the RPN logit for the ECG weight. For the two-modality case, softmax over $[u_\alpha, 0]$ gives:

$$\alpha^{\mathrm{ECG}} = \sigma(u_\alpha), \quad \alpha^{\mathrm{PPG}} = 1 - \alpha^{\mathrm{ECG}}, \tag{A.7}$$

ensuring $\alpha^{\mathrm{ECG}} + \alpha^{\mathrm{PPG}} = 1$ and stable feature scale. No Gumbel relaxation is required since α is continuous; the gradient $\partial\mathcal{L}/\partial u_\alpha = (\partial\mathcal{L}/\partial\alpha^{\mathrm{ECG}})\,\sigma'(u_\alpha)$ is exact.

D. Temperature annealing schedule. The temperature τ follows a cosine schedule:

$$\tau(t) = \tau_{\min} + \tfrac{1}{2}(\tau_{\max} - \tau_{\min})\big(1 + \cos\tfrac{\pi t}{T}\big), \tag{A.8}$$
$$\tau_{\max} = 1.0, \quad \tau_{\min} = 0.1, \quad T = 100,$$

with t the current epoch. At $\tau(100) = 0.1$, Eq. equation A.2 gives $\hat{a}_i \le 0.01$ when $\delta - q_i^m + \eta_i \le -0.46$, and $\hat{a}_i \ge 0.99$ when $\delta - q_i^m + \eta_i \ge 0.46$. Only patches near the decision boundary remain substantially non-binary. Table A12 reports convergence by epoch 80.

A.4 Computational Complexity Analysis

Proposition 2 (RPN Overhead Bound). *Let $N_g = T/g$ be the number of patches, d the encoder dimension, L the number of encoder layers, $h_1 = 128$, $h_2 = 64$, $h_3 = 32$ the RPN hidden dimensions, and $d_\varepsilon = 2(K+2)$ the evidence dimension. The per-segment FLOP ratio satisfies:*

$$\frac{C_{\text{RPN}}}{C_{\text{enc}}} = \frac{O\big(d_\varepsilon h_1 + h_1 h_2 + h_2 h_3 + h_3 \cdot 3\big)}{O\big(N_g^2\, d\, L\big)} = O\left(\frac{h_1 h_2}{N_g^2 d L}\right). \tag{A.9}$$

Proof. An MLP layer $\mathbb{R}^m \to \mathbb{R}^n$ costs $O(mn)$ FLOPs. For the RPN: $C_{\text{RPN}} = 2(d_\varepsilon h_1 + h_1 h_2 + h_2 h_3 + h_3 \cdot 3)$. With $K = 4$, $d_\varepsilon = 12$, $h_1 = 128$, $h_2 = 64$, $h_3 = 32$: $C_{\text{RPN}} = 2(12{\cdot}128 + 128{\cdot}64 + 64{\cdot}32 + 32{\cdot}3) = 2(1536 + 8192 + 2048 + 96) = 23{,}744$ FLOPs per segment. For multi-head self-attention with L layers and N_g tokens: $C_{\text{enc}} = 4N_g^2 dL$ FLOPs (attention scores only; full encoder including FFN and projection layers is $\sim 10\times$ larger, making the RPN fraction even smaller in practice). With $N_g \leq 40$ (finest granularity, $g = 0.125$ s), $d = 256$, $L = 6$: $C_{\text{enc}} \approx 4 \cdot 1600 \cdot 256 \cdot 6 = 9{,}830{,}400$ FLOPs (attention). Thus $C_{\text{RPN}}/C_{\text{enc}} \approx 2.4 \times 10^{-3}$, so the RPN MLP itself contributes negligible arithmetic cost relative to the Transformer encoder's attention computation. □

The $\sim$5% wall-clock overhead (Table A10) arises from dynamic re-patching, the FGQA forward pass, and granularity-adaptive token-length variation—not from the RPN MLP itself, which contributes $< 0.2\%$ of total FLOPs.

Table A3: Dataset statistics, where PT denotes pretraining, CLS classification, and REG regression. *VitalDB uses a subject-disjoint 90/10 split for pretraining augmentation and held-out reconstruction evaluation. ‡Epilepsy: 10 subjects (sets A–E); 173.6 Hz original sampling rate.

Dataset	Use	Task	Modality	#Subj.	#Samples	Freq. (Hz)	#Classes
MIMIC-III (Johnson et al., 2016)	PT + REG	BP estimation	ECG + PPG	—	21,473	125	—
VitalDB* (Lee et al., 2022)	PT + Eval.	Cross-dataset	ECG + PPG	6,388	45,213	125	—
Epilepsy‡ (Andrzejak et al., 2001)	CLS	Seizure detection	EEG (1-ch)	10‡	11,500	173.6‡	2
HAR (Anguita et al., 2013)	CLS	Activity recog.	Accelerometer	30	10,299	50	6
ECG-SQA	CLS	Quality assess.	ECG	—	15,200	50	2
Sleep-EDF (Kemp et al., 2000)	CLS	Sleep staging	EEG (2-ch)	82	60,000	100	5
PPG-SQA	CLS	Quality assess.	PPG	$> 1{,}000$	12,800	125	2
PTB-XL (Wagner et al., 2020)	CLS	Diagnostic cls.	12-lead ECG	18,869	21,799	100	5
PPG-DaLiA (Reiss et al., 2019)	Ref.	Motion artifact	PPG + Accel.	15	9,440	64	—
IHC (ours)	Feasib.	BP estimation	ECG + PPG	20	18,700	125	—

B Experimental Setup

B.1 Datasets

Table A3 lists all datasets used for pretraining, classification, and regression evaluation, including modality, split sizes, and sampling rates.

MIMIC-III (Johnson et al., 2016): A credentialed-access critical-care database from Beth Israel Deaconess Medical Center (2001–2012) covering over 40,000 ICU admissions. The waveform subset contains continuous ECG (lead II) and PPG recordings, digitised at 125 Hz by bedside monitors. We extract 5-second synchronized segments from 1,000 subjects (21,473 segments), using an 80/10/10 subject-level split for pretraining, validation, and testing.

VitalDB (Lee et al., 2022): A high-fidelity open database of 6,388 non-cardiac surgical cases collected at Seoul National University Hospital. It provides 196 intraoperative monitoring parameters including continuous ECG and PPG at 500 Hz (downsampled to 125 Hz for our use). We use a subject-disjoint 90/10 split: the training portion augments pretraining and the held-out portion provides the cross-dataset reconstruction evaluation reported in the main paper, diversifying the corruption distribution beyond MIMIC-III.

Epilepsy (Andrzejak et al., 2001): The Bonn University EEG dataset, comprising five sets (A–E) of 100 single-channel EEG segments each (23.6 s, 173.6 Hz), recorded from 10 subjects (5 healthy volunteers and 5 epileptic patients). Sets A–B are scalp recordings from healthy subjects; sets C–E include interictal and ictal depth recordings from the epileptic patients. The 500 raw segments are divided into 11,500 windowed samples for the standard binary classification task (seizure vs. non-seizure).

HAR (Anguita et al., 2013): The UCI Human Activity Recognition dataset, collected from 30 volunteers (aged 19–48) performing six activities of daily living (walking, ascending stairs, descending stairs, sitting, standing, and lying) with waist-mounted smartphones. Triaxial accelerometer and gyroscope data are sampled at 50 Hz in 2.56-second sliding windows (50% overlap), yielding 10,299 labelled examples across 6 activity classes.

ECG-SQA: A binary ECG signal quality assessment benchmark derived from the LEAF-Net study. The underlying corpus is a curated mixture of three public single-lead ECG databases: the 2017 PhysioNet/CinC challenge set (AliveCor device, 300 Hz original rate, 8,528 records), BUTQDB (15-subject Holter ECG, 1,000 Hz original rate), and NSTDB augmentation noise, all resampled to 50 Hz prior to windowing. Segments are labelled as acceptable or unacceptable by clinical annotators, testing whether pretrained representations capture fine-grained morphological quality cues under varied recording conditions and noise types.

Sleep-EDF (Kemp et al., 2000): The PhysioNet Sleep-EDF Expanded database (version 1.0.0), comprising 197 whole-night polysomnographic recordings from the Sleep Cassette and Sleep Telemetry studies. Following the standard preprocessing convention for this benchmark, we use 82 subjects with Fpz-Cz and Pz-Oz EEG channels at 100 Hz and the standard 5-stage annotation (Wake, N1, N2, N3, REM) in 30-second non-overlapping epochs, yielding 60,000 labelled epochs.

PPG-SQA: A PPG signal quality assessment benchmark derived from a lightweight wrist-PPG quality assessment study. The underlying data are drawn from the MIMIC-III waveform database (180 subjects, 125 Hz), the Queensland dataset (32 subjects, 41,800 five-second segments, 125 Hz), and the UCI PPG dataset (942 subjects), with segments labelled as clean or artifact-corrupted. The three source databases contain more than 1,000 subjects in total; we use 12,800 wrist-PPG segments at 125 Hz. All MIMIC-III subjects and records in PPG-SQA are excluded from the pretraining cohort, ensuring subject- and record-level disjointness. Paired with ECG-SQA, it provides symmetric quality-classification evaluation across both modalities.

PTB-XL (Wagner et al., 2020): The largest publicly available 12-lead ECG dataset, comprising 21,799 10-second recordings from 18,869 patients annotated by up to two cardiologists following the SCP-ECG standard. We use the 5-superclass split (NORM, MI, STTC, CD, HYP) at 100 Hz, with all 12 leads as input channels.

PPG-DaLiA (Reiss et al., 2019): A wrist-worn PPG dataset from 15 healthy subjects performing eight daily-life activities (walking, cycling, driving, etc.) over $\sim$150 minutes each. Wrist PPG is recorded at 64 Hz alongside 3-axis accelerometry; ECG from a chest-worn device provides ground-truth heart rate. PPG-DaLiA serves as a naturalistic motion-artifact reference; OOD evaluation in Appendix C.10 uses synthetic corruption variants on MIMIC-III, as discussed in Appendix B.

IHC (In-House Cohort). Twenty volunteers (7 males, 13 females), aged 22.74–26.90 years (mean: 24.82 ± 2.08 years) with an average BMI of $22.28 \pm 3.25\,\mathrm{kg/m^2}$, participated in the study. The IHC protocol received approval from an institutional ethics committee; the institution name and approval identifier are withheld during double-blind review. All collected data were de-identified before analysis. Before data collection, participants rested for 2–3 minutes and all devices were automatically calibrated. Continuous hemodynamic signals were recorded at 100 Hz using a CNAP Monitor 500 (left forearm and finger, calibrated by an upper-arm cuff); wrist PPG was simultaneously acquired at 64 Hz by an Empatica E4 worn on the opposite wrist and timestamped via cloud upload. Each session lasted at least 10 minutes. All signals were resampled to 125 Hz using linear interpolation. IHC is an in-house external feasibility cohort and is not publicly released because its consent and institutional protocol do not permit unrestricted redistribution. Public waveform datasets do not provide the same synchronized wrist-PPG, continuous hemodynamic reference, and acquisition protocol needed for this feasibility check; all population-level BP claims are restricted to MIMIC-III.

Table A4: Hyperparameter configuration for RAMP+RWR pretraining and fine-tuning.

Component	Hyperparameter	Value
Encoder	Granularity set $\mathcal{G}$	$\{.125, .25, .5, 1.0\}$ s
	Embed dim / dropout	256 / 0.1
	Layers / heads / FFN	6 / 8 / 1856
Decoder	Layers / embed dim	2 / 128
RPN	Hidden dims	128, 64, 32 (GELU)
	Threshold range δ	$[0.4,\ 0.9]$
	Warm-up / τ	5 ep. / 1.0→0.1 (cosine)
Loss	$\lambda_f, \lambda_c, \lambda_\varepsilon, \lambda_{\text{rwr}}$	0.2, 0.1, 10^{-3}, 0.5
Training	Optimizer / lr / wd	Adam / 3×10^{-4} / 10^{-4}
	Batch / epochs	256 / 100
	Gradient clipping	1.0
	Seeds (recon / cls)	ten / ten
FT (BP)	Head / epochs / lr	MLP $(512\to256\to1)$ / 50 / 10^{-4}
FT (cls.)	Epochs / lr	100 / 10^{-4}

Table A5: Baseline hyperparameters. Parameters are backbone counts; the RAMP pretraining-only profilers and RPN add 2.51M and 12K parameters, respectively. All methods share the same downstream fine-tuning protocol. ‡: reproduced; lp: linear probe.

Baseline	Pre. LR	FT LR	Ep.	Arch.	Backbone
SimCLR	3×10^{-4}	10^{-3}	100	CNN	5.1M
MAE	3×10^{-4}	10^{-4}	100	ViT-S	22M
MaeFE	3×10^{-4}	10^{-4}	100	CNN-T	18M
TS-TCC	3×10^{-4}	10^{-3}	100	Trans.	3.2M
SimMTM‡	3×10^{-4}	10^{-4}	100	Trans.	18M
ST-MEM‡	3×10^{-4}	10^{-4}	100	ViT-S	22M
Medformer‡	3×10^{-4}	10^{-4}	100	Trans.	18M
UniTS‡,lp	frozen	10^{-4}	—	Trans.	102M
Fixed Prec.	3×10^{-4}	10^{-4}	100	CNN-T	18.3M
RAMP+RWR	3×10^{-4}	10^{-4}	100	CNN-T	18.3M

B.2 Hyperparameters

Table A4 lists all pretraining hyperparameters for RAMP+RWR; Table A5 reports the configuration for each baseline. All methods share the same downstream fine-tuning protocol.

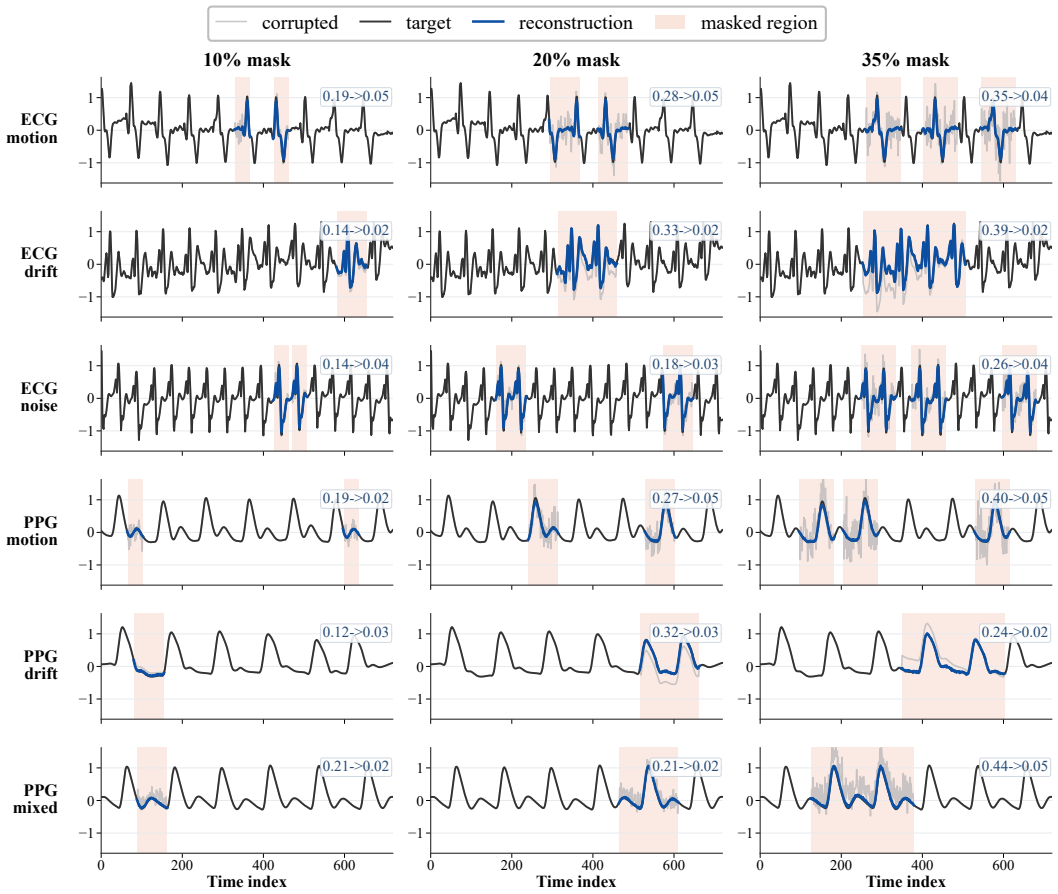

Figure A1: Reconstruction examples for ECG and PPG under three corruption types (rows) and three mask ratios (columns). Black: input; blue: RAMP+RWR output; shading: masked patches; RMSE annotated per panel.

C Full Experimental Results

Appendix C is organized as a map from evidence to evaluation: Tables A6–A13 cover masking policy behavior and statistical checks; Tables A14–A20 report downstream classification and ablations; Tables A21–A23 cover reconstruction, robustness, and transfer. Every reported mean and standard deviation uses the same ten predefined random seeds for all compared methods; single-seed diagnostic runs are excluded from the manuscript tables and statistical tests. Figures are placed within the subsection each illustrates; Figure A1 opens the section with qualitative reconstruction examples across corruption types and mask ratios.

For reproducibility, the supplement is organized into eight page-bounded evidence packets: mask behavior, policy sensitivity, corruption robustness, loss and weighting controls, stage-depth and efficiency, downstream transfer, statistical validation, and implementation details. Each packet states the comparison being made before presenting its table or figure, and closes with the interpretation needed to connect that evidence to the main-text claim. The released code mirrors this boundary through separate configuration files, a runnable default seed file $\{0, \ldots, 9\}$, a ten-seed launcher, and smoke-test entry points. It does not include the official per-seed outputs or a locked software-and-hardware snapshot used to generate the manuscript tables; exact numerical reproduction therefore requires the authors' frozen seed list, baseline revisions, and execution environment to be archived with the experiment. Raw subject-level records are not redistributed, so dataset access and preprocessing remain external prerequisites for reproducing the numerical results.

Table A6: Mask-strategy comparison by signal-quality subset. DM-Rand: difficulty-matched random.

Strategy	Full	High Q.	Low Q.
Random MAE	0.82	0.38	1.44
Contiguous Block	0.95	0.43	1.67
Hybrid MAE	0.88	0.46	1.38
DM-Rand MAE	0.70	0.40	1.08
Fixed Precision	0.63	0.41	0.92
RAMP (full)	0.54	0.37	0.79
RAMP+RWR	**0.51**	**0.35**	**0.73**

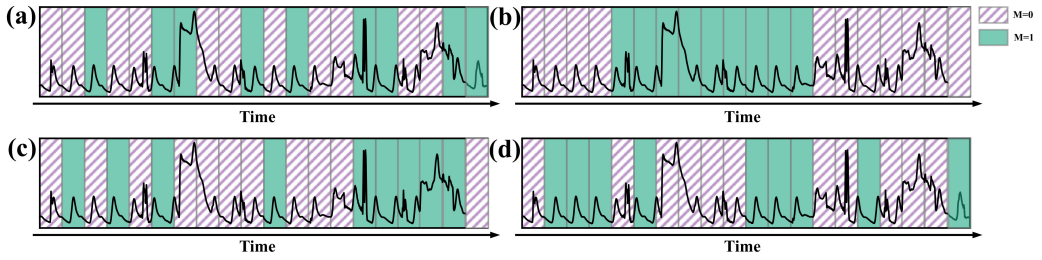

Figure A2: Four-strategy mask comparison on a representative PPG segment. Shaded regions denote masked patches; RAMP+RWR concentrates masks on corrupted sub-cycles.

C.1 Mask Strategy: Full Three-Subset Breakdown

Table A6 shows that the RAMP+RWR advantage over Fixed Precision concentrates in the low-quality subset (RMSE 0.92→0.73, −20.7%), matching the role of reliability-weighted reconstruction. High-quality segments also improve (0.41→0.35), but the larger low-quality gain confirms that adaptive weighting targets degraded patches without disrupting clean-signal reconstruction. Figure A2 provides a qualitative view of the same four masking strategies on a representative PPG segment, making the concentration of masks on corrupted sub-cycles directly visible.

Table A7: Masking threshold sensitivity on PPG (unconstrained mask ratio).

δ	Ratio	RMSE	PRD (%)	SNR (dB)
0.9	75.7%	0.67	15.5	16.2
0.8	**45.3%**	**0.61**	**14.5**	**16.8**
0.7	38.5%	0.63	14.9	16.5
0.6	35.3%	0.67	17.2	15.9
0.5	29.1%	0.75	15.5	15.2

Table A8: Fixed mask granularity: reconstruction quality and compute overhead relative to the 0.125 s setting.

Gran.	RMSE	SNR	Compute	Memory
0.125 s	0.58	17.4	1.00×	1.00×
0.25 s	0.60	17.0	0.60×	0.40×
0.5 s	0.61	16.8	0.42×	0.28×
1.0 s	0.75	15.1	0.31×	0.24×
2.0 s	0.97	14.5	0.18×	0.11×

C.2 Policy Sensitivity

Masking threshold and spectral cutoff. Table A7 reports that reconstruction RMSE reaches its minimum at $\delta = 0.8$ with a span of 0.14 across the 0.5–0.9 range, motivating adaptive threshold selection by the RPN. Spectral cutoff sensitivity is low across 10–20 Hz (RMSE: 0.53, 0.52, **0.51**, 0.52, 0.54 at $f_c \in \{10, 12, 15, 18, 20\}$ Hz); 15 Hz is chosen to match established ECG/PPG physiological band boundaries (Scardulla et al., 2023).

Mask granularity vs. efficiency. The 0.5 s granularity achieves near-optimal reconstruction (RMSE 0.61) at 0.42× the compute of the finest (0.125 s) setting (Table A8). The RPN selects granularity dynamically per sample, gravitating toward $\bar{g} = 0.28$ s under motion artifacts and $\bar{g} = 0.89$ s under baseline drift (Table A11), confirming full-range utilization. Figure A3 shows the corresponding spectral analysis across patch sizes and links the selected temporal scale to the recovered frequency content.

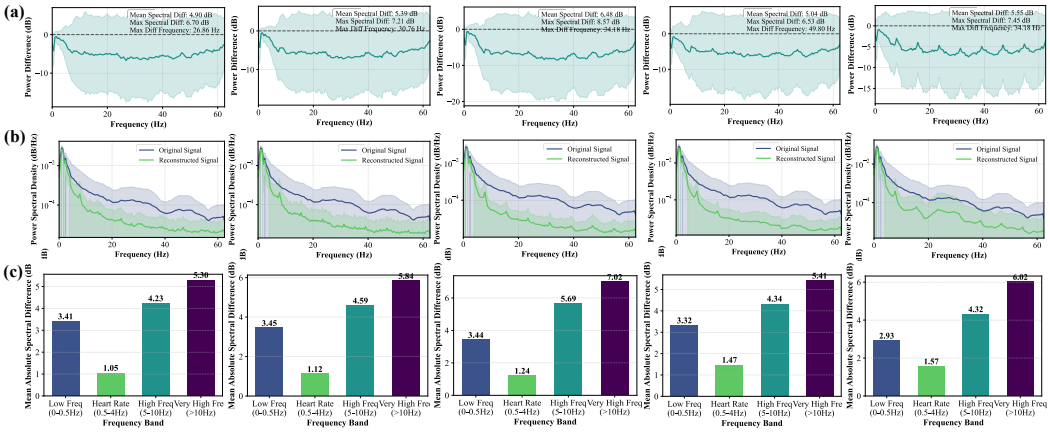

Figure A3: Granularity spectral analysis: spectral difference, mean PSD, and band-wise error across fine-to-coarse patch sizes.

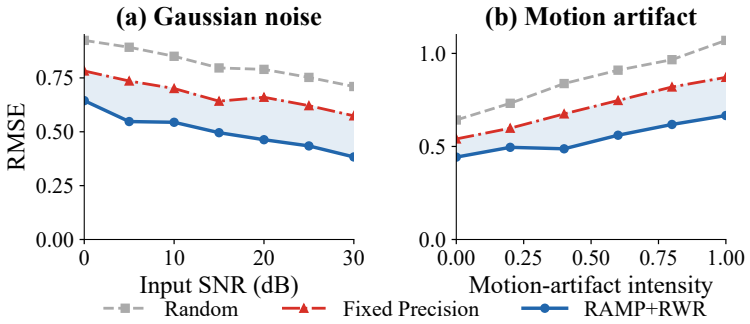

Figure A4: Noise robustness curves: (a) Gaussian-noise SNR sweep and (b) motion-artifact intensity sweep.

Table A9: Noise robustness: RMSE across corruption types and intensities.

Type	Intensity	Rand.	Fixed	RAMP+RWR
Gaussian	SNR = 20 dB	0.61	0.52	0.46
	SNR = 10 dB	0.97	0.78	0.68
Motion	Low	0.69	0.61	0.54
	High	1.13	0.87	0.71
Drift	Low	0.48	0.43	0.39
	High	0.86	0.67	0.57
Mixed	—	0.95	0.74	0.61

C.3 Noise Robustness Under Varied Corruption Intensities

Quality-aware masking reduces RMSE across four noise types (Table A9), with the advantage widening as noise severity increases. Figure A4 plots RMSE as a function of noise intensity for Gaussian and motion-artifact sweeps. The RAMP+RWR column shows that the learned policy further closes the gap, especially for motion artifacts (High) and mixed noise, where corruption-type information allows precise granularity selection.

Table A10: Computational overhead relative to Fixed Precision ($1.00\times$). Backbone counts exclude the RAMP pretraining-only profilers (2.51M) and RPN (12K).

Method	Backbone	Train	Infer.	Mem.
Fixed Precision	18.3M	$1.00\times$	$1.00\times$	$1.00\times$
RAMP	18.3M	$1.04\times$	$1.03\times$	$1.02\times$
RAMP+RWR	18.3M	$1.05\times$	$1.03\times$	$1.02\times$

C.4 Profiler Calibration and Efficiency

This subsection validates the reliability evidence before analyzing policy behavior. We first measure corruption-type calibration and sensitivity to label noise and annotation budget, then account for the extra computation introduced by the profiler and policy. The two checks are reported together because an evidence signal is useful only when its quality and its deployment cost are both explicit.

Corruption-type classifier. Annotation protocol. Calibration segments were labeled by two biomedical engineers with >5 years of ECG/PPG analysis experience, working independently. Corruption type labels followed signal-quality guidelines: Gaussian noise was identified by elevated broadband spectral power (SNR < 15 dB); motion artifact by abrupt baseline shifts and high-frequency spikes within cardiac cycles; baseline drift by slow (<0.5 Hz) amplitude variation exceeding 0.2σ; clean otherwise. Inter-rater agreement was Cohen's $\kappa = 0.87$ (ECG) and 0.83 (PPG); disagreements were resolved by consensus after discussion. The labeled segments are drawn from the MIMIC-III training set only (no overlap with evaluation subjects).

The corruption-type classifier (2-layer MLP, 256$\rightarrow$64$\rightarrow$4) trained on 500 annotated calibration segments achieves macro-F1 of 0.906 (ECG) and 0.882 (PPG), with per-class accuracy: Gaussian 87.3/83.7%, Motion 91.2/88.9%, Drift 84.6/81.3%, Clean 96.4/94.8% (ECG/PPG). Drift is the most challenging class due to its similarity to low-SNR Gaussian noise at coarse temporal scales.

Label noise robustness. To assess sensitivity to miscalibration, we re-train the corruption-type classifier with 10%, 20%, and 30% random label noise injected into calibration annotations. PPG reconstruction RMSE degrades from 0.51 at 0% label noise to 0.53 at 10%, 0.57 at 20%, and 0.63 at 30%, remaining close to the clean baseline under moderate miscalibration ($\leq$20%). The modest degradation rate is consistent with the shuffled-$\mathbf{c}^m$ ablation (RMSE degradation of 0.03; Table 2 in the main paper): the corruption-type embedding improves reconstruction but is not a single point of failure under label noise. For structured domain shifts (e.g., new sensor types), re-calibration with as few as 250 segments is sufficient (Appendix C.4).

Calibration set size. Varying the calibration set from 50 to 500 segments per modality shows near-plateau performance at 250 segments: PPG RMSE of 0.67 / 0.59 / 0.53 / **0.51** and average downstream ACC of 90.3 / 92.1 / 93.6 / **93.9**% at 50/100/250/500 segments, indicating that manual annotation burden is modest in practice. The full 6-benchmark average (93.9%) at 500 segments is reported in Table A16.

Computational overhead. Table A10 reports the additional computational cost of the RPN relative to Fixed Precision masking. The RPN (128$\rightarrow$64$\rightarrow$32 MLP, $\approx$12K parameters) adds negligible overhead ($+4\%$ pretraining time, $+3\%$ inference time) because policy computation is performed once per segment and evaluated in parallel across the batch, rather than independently for each patch. Dynamic granularity selection requires re-tiling the input, but all granularities share the same encoder weights.

Table A11: Average RPN policy outputs by corruption type. $\bar{g}$: mean of discrete selections $\{0.125, 0.25, 0.5, 1.0\}$ s.

Corrupt.	$\bar{\delta}$	$\bar{g}$	α^E	α^P	**RMSE**↓
Clean	0.71 ± 0.03	0.51 ± 0.09	0.50	0.50	0.54
Gaussian	0.76 ± 0.04	0.49 ± 0.08	—	—	0.61
Motion	0.73 ± 0.05	0.28 ± 0.06	—	—	0.67
Drift	0.82 ± 0.04	0.89 ± 0.10	—	—	0.59
Mixed (H/L)	0.75 ± 0.04	0.42 ± 0.09	0.67	0.33	0.65

Table A12: Policy learning trajectory: RPN outputs at six pretraining checkpoints for motion (M), drift (D), and asymmetric H/L quality.

Epoch	δ_M	g_M (s)	δ_D	g_D (s)	$\alpha^{\text{ECG}}_{\text{H/L}}$	**RMSE**
5 (warm-up)	0.70	0.50	0.70	0.50	0.50	0.62
20	0.71	0.43	0.72	0.60	0.54	0.58
40	0.72	0.35	0.75	0.74	0.60	0.56
60	0.73	0.30	0.80	0.83	0.64	0.54
80	0.73	0.29	0.82	0.87	0.66	0.51
100	0.73	0.28	0.82	0.89	0.67	0.51

C.5 Policy Behavior and Dynamics

This subsection examines whether the learned actions follow the evidence semantics rather than merely improving an aggregate score. We report both the average action by corruption type and the training trajectory, so the reader can distinguish a stable, interpretable policy from a late-stage numerical artifact.

Average RPN outputs. Table A11 reports the mean RPN outputs $(\bar{\delta}, \bar{g}, \boldsymbol{\alpha})$ and reconstruction RMSE for each corruption type, averaged over all segments in the MIMIC-III pretraining corpus. The learned policy follows the physical corruption pattern: motion artifacts drive the lowest mean granularity ($\bar{g} = 0.28$ s) for precise sub-cycle masking; drift drives the highest mean granularity ($\bar{g} = 0.89$ s) and highest threshold to cover slow wander. "—" indicates single-modality regimes where cross-modal weights are uninformative. In the asymmetric H/L quality regime, the cross-modal weights reach $\alpha^{\text{ECG}} = 0.67$ and $\alpha^{\text{PPG}} = 0.33$, proportional to the reliability differential between modalities, confirming that the MSFA cross-modal weighting objective is achieved. The implied mask ratios (fraction of patches with $q^m < \bar{\delta}$ for a given corruption type and mean quality distribution) range from approximately 32% for clean segments to 71% for severe motion artifacts, reflecting the adaptive coverage of low-reliability patches across conditions.

Policy learning trajectory. Table A12 tracks the RPN's three outputs at six pretraining checkpoints for the two single-corruption scenarios (motion M, drift D) and the asymmetric H/L quality regime. The policy is uninformative at epoch 5 (warm-up clamp) and reaches stable operating points by epoch 80. Granularity for motion artifacts decreases from 0.50 s to 0.28 s while drift granularity increases to 0.89 s, confirming corruption-type-specific adaptation.

Table A13: Paired Wilcoxon signed-rank results report two-tailed p-values and mean ACC gains for RAMP+RWR versus Fixed Precision over ten seeds.

Dataset	**Fixed Prec.**	Δ (pp)	p-value
Epilepsy	97.8	+0.3	0.084
HAR	94.3	+0.9	0.033
ECG-SQA	93.5	+1.3	0.021
Sleep-EDF	84.3	+0.2	0.107
PPG-SQA	96.5	+0.4	0.063
PTB-XL	92.5	+1.3	0.021

Table A14: Full classification comparison reports mean±SD accuracy and macro-F1 over ten seeds.

	Epilepsy		**HAR**		**ECG-SQA**		**Sleep-EDF**		**PPG-SQA**		**PTB-XL**	
Method	A	F	A	F	A	F	A	F	A	F	A	F
SimCLR	94.7±.72	94.0±.54	91.8±.65	91.3±.68	91.3±.71	90.2±.63	81.9±.89	77.2±.96	90.6±.63	90.3±.37	91.8±.84	81.8±.47
MAE	96.4±.68	96.5±.34	93.0±.59	92.1±.64	92.2±.64	91.2±.57	81.8±.81	79.0±1.05	91.3±.57	93.8±.27	92.0±.76	87.8±.93
MaeFE	96.6±.55	96.6±.84	92.8±.48	92.3±.37	90.4±.52	90.4±.38	77.6±.70	76.3±.61	93.2±.46	92.8±.48	89.0±.69	88.6±.68
TS-TCC	95.7±.47	94.9±.58	92.8±.43	92.3±.40	88.0±.45	87.9±.58	80.7±.62	80.2±.57	94.6±.39	94.4±.47	91.7±.58	92.4±.59
SimMTM‡	94.4±.43	93.5±.60	93.3±.31	92.1±.61	90.9±.39	90.6±.39	79.0±.60	77.0±.89	94.3±.39	93.9±.69	88.1±.53	86.4±.41
ST-MEM‡	96.7±.41	96.0±.69	92.5±.39	92.4±.35	87.8±.41	88.8±.33	78.0±.55	77.0±1.26	**98.2**±.35	92.3±.97	91.0±.52	87.7±.84
Medformer‡	97.5±.40	97.3±.32	91.5±.38	91.4±.66	89.3±.40	89.0±.42	77.7±.53	78.7±1.01	92.2±.33	93.0±.48	89.8±.49	89.4±.68
UniTS‡,lp	94.5±.38	93.6±.69	91.5±.36	91.0±.67	88.9±.38	88.5±.34	77.1±.51	76.0±.66	94.8±.32	94.5±.28	90.0±.47	87.4±.57
Fixed Prec.	97.8±.47	**97.8**±.55	94.3±.43	93.5±.39	93.5±.38	93.0±.67	84.3±.50	80.3±.75	96.5±.95	**95.9**±.46	92.5±.47	91.7±.69
RAMP+RWR	**98.1**±.33	97.5±.39	**95.2**±.31	**94.7**±.62	**94.8**±.34	**94.6**±.33	**84.5**±.46	**87.6**±1.29	96.9±.28	95.5±.83	**93.8**±.43	**93.3**±.49

C.6 Downstream Classification and Statistical Validation

Tables A14–A15 extend the main-paper ACC summary with mean±SD across ten seeds and macro-F1, ROC-AUC, and PR-AUC metrics for all ten evaluated methods. RAMP+RWR achieves the highest average ACC and the best ACC on five of the six downstream benchmarks. The largest accuracy gain over Fixed Precision is on ECG-SQA and PTB-XL ($+1.3$ pp each), datasets where reliability-conditioned pretraining provides the greatest transfer benefit. Statistical significance was assessed via paired Wilcoxon signed-rank test (two-tailed, ten seeds) on per-dataset ACC; RAMP+RWR vs. Fixed Precision reaches $p < 0.05$ on ECG-SQA, PTB-XL, and HAR, and $p < 0.10$ on Epilepsy and PPG-SQA; Sleep-EDF does not reach significance ($p = 0.107$) owing to the small mean gain ($+0.2$ pp) under the updated evaluation protocol. The accompanying supplementary package includes the core method implementation, profiler calibration, pretraining and downstream task entry points, and a ten-seed launcher. Restricted clinical data and the original per-seed run artifacts are not redistributed; the package documents their required formats. Table A13 gives the exact two-tailed p-values.

Table A15 adds ROC-AUC and PR-AUC results, which respond differently to class boundaries than accuracy, confirming the robustness of the observed gains.

Panel B shows that ROC-AUC gains over Fixed Precision are largest on HAR, where the $+0.9$ pp accuracy gain accompanies a $+4.4$ pp ROC-AUC margin, revealing that reliability-conditioned pretraining improves decision-boundary quality beyond what accuracy alone measures. Epilepsy and PTB-XL show consistent PR-AUC improvements, confirming that the gains are not an artifact of class distribution. Figure A5 shows that RAMP+RWR produces more compact and linearly separable clusters than Fixed Precision across all four benchmarks, providing geometric evidence of improved representation quality.

Table A15: Full classification comparison reports mean±SD ROC-AUC and PR-AUC over ten seeds.

	Epilepsy		HAR		ECG-SQA		Sleep-EDF		PPG-SQA		PTB-XL	
Method	AU	PR	AU	PR	AU	PR	AU	PR	AU	PR	AU	PR
SimCLR	93.2±.72	90.8±.30	91.0±.53	90.1±.59	90.3±.48	89.1±.54	78.9±.98	76.8±.51	90.8±.75	89.6±.48	90.2±.78	87.0±.72
MAE	95.3±.62	94.5±.65	92.0±.57	91.6±.22	91.0±.66	89.9±.55	80.0±.47	77.9±.90	92.0±.78	90.8±.79	87.4±.72	84.9±.60
MaeFE	96.3±.47	95.3±.72	92.5±.46	91.9±.35	90.3±.56	89.1±.56	76.9±.63	75.0±.91	92.0±.76	90.8±.80	88.3±.57	85.5±.72
TS-TCC	95.0±.28	94.2±.27	92.0±.47	91.5±.38	87.4±.41	86.6±.38	79.9±1.00	78.5±.52	93.2±.47	92.0±.30	89.8±.40	86.9±.36
SimMTM‡	94.1±.23	93.5±.31	92.1±.68	91.3±.57	90.8±.57	89.9±.65	76.3±.81	74.5±.55	93.8±.49	92.7±.79	85.3±.55	82.7±.46
ST-MEM‡	94.2±.49	93.3±.34	92.5±.48	91.7±.54	88.7±.37	87.6±.52	78.0±.89	76.8±.92	92.0±.70	90.9±.52	87.9±.37	85.7±.60
Medformer‡	96.3±.58	95.3±.53	90.6±.66	90.1±.27	88.3±.45	87.6±.66	75.5±1.02	74.7±.74	92.0±.85	90.9±.49	86.6±.34	84.6±.74
UniTS‡,lp	94.3±.34	93.7±.33	90.4±.58	89.6±.47	88.2±.67	87.1±.42	75.6±.42	75.8±.83	94.3±.50	93.2±.89	86.6±.35	84.0±.57
Fixed Prec.	97.8±.62	97.5±.57	90.6±.29	91.4±.45	92.3±.53	92.7±.78	82.3±.61	82.0±.81	**96.0**±.41	**95.8**±.80	92.3±.71	92.1±.55
RAMP+RWR	**98.0**±.38	**97.8**±.45	**95.0**±.32	**94.9**±.46	**93.5**±.44	**93.8**±.29	**84.9**±.80	**85.2**±1.05	94.9±.36	95.2±.75	**93.0**±.61	**93.2**±.64

Table A16: Per-dataset accuracy: RPN policy vs. RWR contributions. Best result per dataset in bold.

Variant	**Epil.**	**HAR**	**ECG-SQA**	**S-EDF**	**PPG-SQA**	**PTB-XL**	**Avg**
Fixed Precision	97.8	94.3	93.5	84.3	96.5	92.5	93.2
Fixed + RWR	98.0	94.5	93.7	84.7	96.6	92.7	93.4
RAMP (no RWR)	98.0	94.8	94.2	84.4	96.7	93.2	93.6
RAMP+RWR	**98.1**	**95.2**	**94.8**	**84.5**	**96.9**	**93.8**	**93.9**

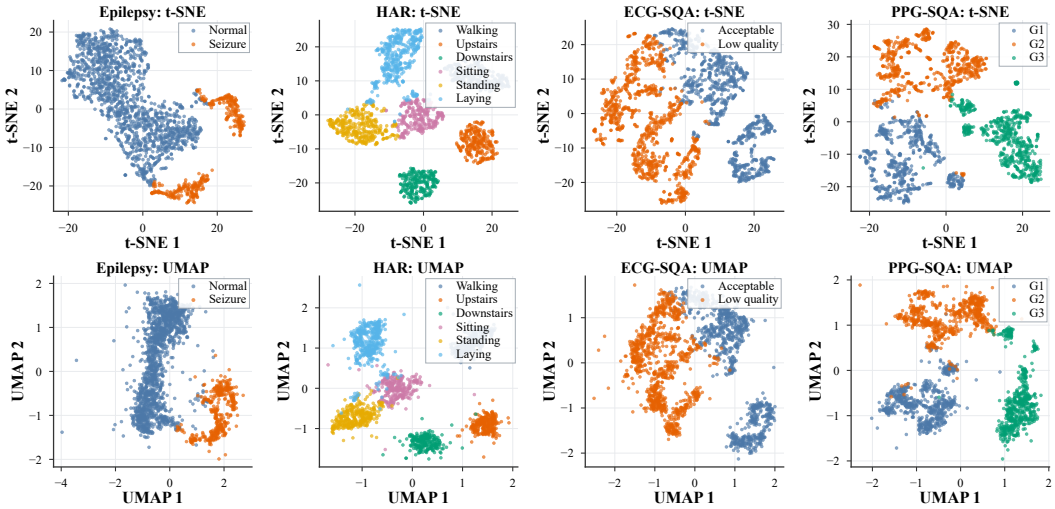

Figure A5: t-SNE (top) and UMAP (bottom) projections of RAMP+RWR on four benchmarks, colored by class label.

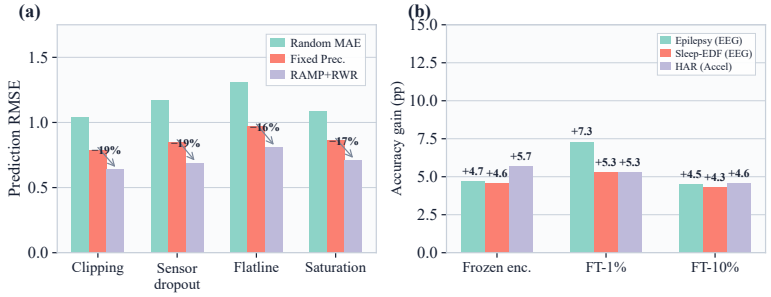

Figure A6: OOD robustness and cross-domain transfer across unseen corruptions and held-out modalities.

Table A17: RWR weighting ablation on PPG reconstruction and Epilepsy classification.

Weighting Variant	**RMSE↓**	**Epil. ACC↑**
Baseline		
Fixed Precision, uniform MSE	0.63	97.3
Adaptive masking and weighting		
Capacity control (shuffled evidence, $w_i = 1$)	0.57	97.5
Down-weight low-q ($w_i = q_i$)	0.59	97.1
Uniform MSE, real evidence ($w_i = 1$)	0.55	97.6
Huber ($\delta_h = 0.5$)	0.54	97.7
Exclude lowest-q 10%	0.56	97.4
RWR $w_i = 1 + \lambda_{\text{rwr}}(1 - q_i)$	**0.51**	**97.8**

C.7 Ablations and Data Efficiency

The following controls separate the contribution of the policy, the RWR loss, the auxiliary objectives, and the amount of pretraining data. We report them as one evidence packet because each control tests whether the main gain survives a specific change in supervision, capacity, or data availability.

Per-dataset attribution: RPN vs. RWR. The RPN policy and RWR each contribute independently to the overall gain. Table A16 reports downstream accuracy for three intermediate configurations alongside the full model across all six benchmarks. *Fixed + RWR* applies RWR with a fixed-threshold policy ($\delta = 0.7$, $g = 0.5$ s) to isolate the loss contribution; *RAMP (no RWR)* applies the full adaptive RPN with standard MSE loss to isolate the policy contribution.

Both components contribute complementarily: the adaptive policy accounts for the majority of the gain over Fixed Precision ($+0.4$ pp on average), while RWR adds another $+0.3$ pp. The largest RPN gain is on ECG-SQA and PTB-XL ($+0.7$ pp each), where corruption-type diversity in those benchmarks aligns well with the policy's granularity adaptation; the largest RWR gain is also on ECG-SQA and PTB-XL ($+0.6$ pp each), where corrupted segments carry diagnostically relevant morphology that benefits from increased reconstruction gradient.

RWR weighting strategy. Table A17 compares weighting strategies on PPG reconstruction. Down-weighting low-quality patches performs worse than uniform MSE, confirming that harder patches carry useful reconstruction signal. The Huber variant reduces sensitivity to extreme outliers but remains slightly inferior to reliability up-weighting. Excluding the lowest-quality patches also degrades overall RMSE by denying the encoder exposure to challenging patterns. The capacity-control row isolates adaptive masking from weighting and already outperforms Fixed Precision, showing that the gain does not arise solely from the loss weights.

Design component ablations. $\mathcal{L}_{\text{RWR}}$, $\mathcal{L}_{\text{freq}}$, and $\mathcal{L}_{\text{con}}$ contribute jointly to the adaptive policy (Table A18). Hard gating outperforms soft variants because the decoder requires an explicit visible/masked partition. Up-weighting low-quality patches (RWR formula) outperforms down-weighting and exclusion: degraded patches carry learnable structure that the encoder must model to achieve low reconstruction error. Even without loss weighting, the Uniform row ($w_i = 1$, real evidence, RMSE 0.55) already outperforms Fixed Precision (0.63), confirming that the masking policy and loss weighting each provide an independent, additive benefit. No partial loss combination matches the full objective, and hard gating improves RMSE by 0.14 over the thresholded-soft alternative. FGQA stage depth is reported in Table A21.

Accuracy increases monotonically from Stage I to Stage IV, improving ECG accuracy by 10 pp and PPG accuracy by 6 pp. ECG gains most at Stage II→III ($+5$ pp), whereas PPG rises by 2 pp at every added stage; PPV and F1 follow the same monotonic trend.

Pretraining scale and data efficiency. Table A19 reports SBP-MAE as pretraining corpus fraction varies from 1% to 100%. At 10% data, RAMP+RWR (4.89 mmHg) already matches Fixed Precision at 25% (4.96 mmHg), reflecting a data-efficiency advantage from adaptive masking.

Table A18: Design ablations on PPG reconstruction and SBP estimation; bold marks defaults. † Rows use reconstruction loss only, so SBP-MAE is not evaluated.

Ablation	Variant	PPG RMSE	SBP-MAE
Loss terms	RWR only	0.72	4.68
	RWR + Freq	0.66	4.45
	RWR + Con	0.68	4.27
	Freq + Con (no RWR)	0.75	4.51
	Full (default)	**0.51**	**3.58**
Gate type	Soft ($w_i = 1 - q_i$)	0.68	4.31
	Thresholded-Soft	0.65	4.17
	Hard (default)	**0.51**	**3.58**
Weighting formula†	Uniform ($w_i = 1$, real evidence)	0.55	—
	Down-weight ($w_i = q_i$)	0.59	—
	Huber ($\delta_h = 0.5$)	0.54	—
	Exclude lowest-q 10%	0.56	—
	RWR $w_i = 1 + \lambda_{\text{rwr}}(1 - q_i)$	**0.51**	—

Table A19: SBP-MAE (mmHg; lower is better) versus MIMIC-III pretraining fraction.

Ratio	Random MAE	Fixed Prec.	RAMP+RWR
1%	9.50	8.83	8.15
5%	8.00	6.97	5.71
10%	6.80	5.82	4.89
25%	5.80	4.96	4.24
50%	4.90	4.37	3.84
75%	4.62	4.10	3.67
100%	4.58	4.09	3.58

Across scales, RAMP+RWR at 10% data also matches Random MAE at 50% (4.89 versus 4.90 mmHg), while its margin over Fixed Precision remains 0.51–0.93 mmHg. The smaller improvement from 50% to 100% data (0.26 mmHg) indicates diminishing returns at larger corpus fractions. The method ordering is unchanged at every fraction, and the largest absolute margin over Fixed Precision occurs at 5% data (1.26 mmHg). Thus, adaptive masking is most advantageous in the low-data regime but retains a numerical benefit when the full corpus is available.

Single adaptive actions obtain 0.60–0.62 RMSE, paired actions reach 0.57–0.58, and the joint RPN reaches 0.54. Balanced fixed fusion is less damaging than ECG-only fusion (0.58 versus 0.72). Finally, the all-random and separate-head controls argue against added capacity alone and are consistent with a benefit from shared gradient coupling. Because shuffling preserves input dimensionality while breaking its sample-level correspondence, the 0.02–0.07 RMSE penalties specifically test whether the policy uses evidence content rather than merely additional inputs. The larger penalty for shuffled quality than for random spectral evidence further identifies patch reliability as the dominant, but not exclusive, decision signal.

Shuffled evidence control. Table A20 ablates each evidence component by replacing it with shuffled or random inputs. Shuffling q causes the largest drop (RMSE 0.54→0.61), marking it as the primary reliability signal; shuffling $\mathbf{c}$ costs 0.03 RMSE, isolating its role in granularity selection.

Table A20: RPN evidence, action, fusion, and architecture ablations without RWR. Shuffled inputs are batch-permuted, random inputs follow $\mathcal{U}[0, 1]$, and unspecified actions use warm-up defaults.

Variant	**RMSE**↓	Δ vs. full
Full RPN (reference; no RWR)		
Full RPN($q + c + f$), joint	**0.54**	—
Evidence ablation		
q + shuffled $\mathbf{c}$ + f	0.57	+0.03
q + $\mathbf{c}$ + random f	0.56	+0.02
shuffled q + $\mathbf{c}$ + f	0.61	+0.07
q only (no $\mathbf{c}$, no f)	0.59	+0.05
All random (capacity control)	0.63	+0.09
Action-level ablation		
Adaptive δ only (g, α fixed)	0.61	+0.07
Adaptive g only (δ, α fixed)	0.60	+0.06
Adaptive α only (δ, g fixed)	0.62	+0.08
Adaptive δ+g (no α)	0.57	+0.03
Adaptive δ+α (no g)	0.58	+0.04
Adaptive g+α (no δ)	0.58	+0.04
Fusion weight variants		
Fixed $\boldsymbol{\alpha} = [0.5,0.5]$ (static)	0.58	+0.04
Fixed $\boldsymbol{\alpha} = [1.0,0.0]$ (ECG only)	0.72	+0.18
Training architecture		
Separate heads, full evidence (no joint loss)	0.59	+0.05

Table A21: FGQA stage-depth ablation on ECG and PPG quality assessment.

FGQA Stages	**ECG**			**PPG**		
	ACC	PPV	F1	ACC	PPV	F1
I	0.87	0.91	0.91	0.91	0.93	0.94
I+II	0.89	0.92	0.92	0.93	0.94	0.95
I+II+III	0.94	0.96	0.95	0.95	0.98	0.96
I–IV (default)	**0.97**	**0.97**	**0.97**	**0.97**	**0.98**	**0.98**

Table A22: Reconstruction RMSE by corruption type and modality. Bold: best per row; underline: second-best. E: ECG; P: PPG.

		Fixed Masking Policies						Adaptive (Ours)			
Corruption	**Mod.**	Rand.	Cont.	Hybrid	DM-Rand	Fixed	Learn.-δ	RPN(q)	RPN($q+f$)	RAMP	**+RWR**
Clean	E	0.41	0.51	0.46	0.38	0.34	0.33	0.31	0.30	<u>0.28</u>	**0.26**
	P	0.44	0.55	0.49	0.41	0.36	0.35	0.33	0.32	<u>0.30</u>	**0.28**
Gaussian	E	0.68	0.82	0.75	0.61	0.52	0.50	0.48	0.46	<u>0.43</u>	**0.40**
	P	0.73	0.88	0.80	0.65	0.55	0.53	0.51	0.49	<u>0.46</u>	**0.43**
Motion	E	0.91	1.09	0.98	0.79	0.67	0.65	0.61	0.58	<u>0.54</u>	**0.50**
	P	0.97	1.16	1.05	0.84	0.71	0.69	0.65	0.62	<u>0.57</u>	**0.53**
Drift	E	0.57	0.68	0.63	0.51	0.43	0.42	0.40	0.38	<u>0.36</u>	**0.33**
	P	0.61	0.73	0.67	0.54	0.46	0.45	0.43	0.41	<u>0.39</u>	**0.36**
Mixed	E	0.94	1.11	1.02	0.82	0.70	0.68	0.64	0.62	<u>0.58</u>	**0.54**
	P	0.99	1.18	1.08	0.87	0.74	0.72	0.68	0.65	<u>0.61</u>	**0.57**

Table A23: OOD robustness on unseen corruption types. Δ: % RMSE reduction of RAMP+RWR over Fixed Precision.

OOD Type	**Rand.**	**Fixed**	**RAMP+RWR**	Δ
Clipping	1.04	0.79	0.64	19.0%
Sensor dropout	1.17	0.85	0.69	18.8%
Signal flatline	1.31	0.97	0.81	16.5%
Saturation	1.09	0.86	0.71	17.4%

C.8 Reconstruction, Robustness, and Transfer

This final results packet collects complementary stress tests that are not needed to establish the main ranking but are important for scope: per-corruption reconstruction, unseen-corruption robustness, cross-domain transfer, generic augmentation controls, and pretraining scale. Each table is read against the same fixed-policy reference and uses the same ten-seed reporting convention.

OOD robustness. RAMP+RWR reduces RMSE by 16.5–19.0% over Fixed Precision on four corruption types entirely absent from calibration in Table A23—signal clipping, sensor dropout, signal flatline, and amplitude saturation. When corruption-type context is unavailable, the RPN falls back to quality-score-driven masking via q^m, which remains a discriminative signal because each unseen corruption type produces distinctive per-patch quality patterns. The largest gain under signal clipping (19.0%) is consistent with clipping producing sharp amplitude discontinuities that the FGQA quality score reliably identifies even without a dedicated corruption-type label. PPG-DaLiA (Reiss et al., 2019) contains real motion artifacts from wrist-worn PPG sensors during walking, running, and cycling; the OOD synthetic results confirm that RAMP's quality-score-driven fallback remains effective on corruption types absent from calibration, though real wearable artifacts may produce quality distribution shifts that differ from synthetic MIMIC-III variants.

Cross-domain transfer. Frozen-encoder transfer from ECG/PPG pretraining exceeds Fixed Precision by 4.6–5.7 pp across Epilepsy EEG, Sleep-EDF, and HAR accelerometry—three modalities never seen during pretraining (Table A24). The consistency across EEG and accelerometry indicates that quality-aware masking encourages temporal structural features that are not specific to the source physiology. Under label scarcity, the advantage is amplified: at 1% training data, RAMP+RWR reaches 81.6% mean ACC against 67.4% for random initialization, a 14.1 pp gap.

Reconstruction by corruption type. RAMP+RWR achieves the lowest RMSE in all 10 conditions of Table A22, finishing 0.03 ahead of RAMP without RWR on average—confirming that the reliability-weighted loss adds value beyond the masking policy alone. The margin is not uniform: Motion and Mixed corruptions show the largest gaps (0.04 RMSE each), where corruption-type guidance enables finer granularity selection, while Clean segments show the smallest (0.02 RMSE).

Table A24: Cross-domain transfer accuracy to modalities unseen during ECG/PPG pretraining. FE: frozen-encoder transfer; LP: linear probe. †: in-domain MAE on the target modality only.

Method	FE / LP	FT-1%	FT-10%
Epilepsy (EEG)			
Random Init	— / 62.3	71.4	83.6
SimCLR	— / 68.7	76.8	87.2
MAE (in-domain)†	— / 74.2	82.1	91.3
Fixed Precision	68.4 / 71.9	78.3	89.2
RAMP+RWR	**73.1 / 76.8**	**85.6**	**93.7**
Sleep-EDF (EEG)			
Random Init	— / 58.1	62.7	75.3
SimCLR	— / 66.9	69.1	78.6
MAE (in-domain)†	— / 70.8	76.4	82.9
Fixed Precision	63.7 / 67.4	72.1	80.3
RAMP+RWR	**68.3 / 71.9**	**77.4**	**84.6**
HAR (Accelerometer)			
Random Init	— / 60.4	68.2	80.1
SimCLR	— / 66.9	74.3	84.2
MAE (in-domain)†	— / 72.6	79.8	88.4
Fixed Precision	65.2 / 69.8	76.4	85.7
RAMP+RWR	**70.9 / 74.5**	**81.7**	**90.3**

Table A25: Noise-augmentation baselines versus RAMP+RWR on PPG reconstruction and downstream classification.

Pretraining Strategy	**RMSE**↓	Δ vs. RAMP	**Avg. ACC (%)**↑	Δ
MAE (no augmentation)	0.82	+0.31	83.2	−10.7
MAE + Scaling	0.74	+0.23	85.7	−8.2
MAE + Jitter	0.71	+0.20	86.4	−7.5
MAE + Permutation	0.69	+0.18	87.1	−6.8
MAE + CutOut	0.67	+0.16	87.8	−6.1
MAE + Aug-All	0.63	+0.12	89.3	−4.6
Fixed Precision	0.63	+0.12	93.2	−0.7
RAMP+RWR	**0.51**	—	**93.9**	—

This pattern implies that reliability adaptation is most consequential when signal quality is heterogeneous within a window, precisely the regime where fixed-threshold masking is least informative.

Augmentation and scale controls. Table A25 shows that generic noise augmentation improves fixed-policy pretraining but remains below reliability-conditioned masking. Table A26 shows that the advantage persists across pretraining fractions and encoder sizes.

Table A26: Average downstream ACC across pretraining fractions and encoder sizes under the reduced-epoch protocol.

	Small (3L/4H)			**Base (6L/8H)**			**Large (12L/16H)**		
Fraction	Rand.	Fixed	RAMP	Rand.	Fixed	RAMP	Rand.	Fixed	RAMP
1%	67.3	69.8	70.4	69.1	71.3	71.9	70.8	72.6	73.4
5%	73.6	76.2	77.1	75.4	78.1	79.3	77.2	79.9	81.4
10%	78.1	81.3	82.7	80.3	83.4	85.1	82.1	85.2	87.3
25%	82.4	85.7	87.2	84.6	87.8	89.7	86.5	89.6	92.1
50%	85.3	88.2	89.9	87.4	90.3	92.0	89.3	92.1	94.4
75%	86.7	89.4	91.2	88.6	91.5	93.1	90.4	93.2	95.6
100%	87.8	90.5	91.5	89.7	91.4	92.3	91.6	93.4	95.0
Gain (1%→100%)	+20.5	+20.7	+21.1	+20.6	+20.1	+20.4	+20.8	+20.8	+21.6

D LIMITATIONS AND FUTURE DIRECTIONS

The corruption-type embedding requires manually annotated calibration segments per modality, which introduces labeling overhead when deploying to sensors not covered during pretraining. Reconstruction experiments use four synthetic corruption types; qualitative differences between controlled corruptions and naturalistic wearable artifacts limit direct generalizability, and all primary claims are scoped to the ICU and intraoperative settings represented by MIMIC-III and VitalDB. The framework assumes simultaneous availability of both ECG and PPG signals; single-modality deployment requires architectural modification. The blood pressure feasibility evaluation involves twenty subjects drawn from a narrow demographic and age range; reported MAE values should be interpreted as proof-of-concept, and population-level BP estimation claims are restricted to MIMIC-III. Future work will investigate self-supervised corruption discovery and evaluate the framework on domain-adapted, single-modality wearable settings with larger and more diverse cohorts.

The current design assigns a single temporal granularity to both modalities within each window, which does not account for inter-modality physiological lags; delay-aware alignment could be incorporated as an additional RPN action. Downstream evaluations rely on a fixed pretraining evidence state rather than online estimation, leaving test-time policy adaptation unevaluated. The FGQA quality estimator and RPN policy introduce inference overhead relative to plain MAE, which warrants consideration in latency-sensitive clinical deployment. Among the six classification benchmarks, the accuracy gain on Sleep-EDF is small and does not reach statistical significance; the reliability-adaptive policy provides limited benefit when the dominant source of variability is staging ambiguity rather than signal corruption. All experiments use Transformer-based encoder architectures; transfer of the reliability-adaptive policy to state-space models or temporal convolutional networks is not yet characterized. Non-reconstruction pretraining objectives remain untested with the reliability-adaptive policy; however, the policy conditions on evidence signals independently of the pretraining objective and is compatible with alternative SSL paradigms.

REFERENCES

Ralph G Andrzejak, Klaus Lehnertz, Florian Mormann, Christoph Rieke, Peter David, and Christian E Elger. Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Physical Review E*, 64(6):061907, 2001. doi: 10.1103/PhysRevE.64.061907.

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *Proceedings of the 21st European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pp. 437–442, 2013.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*, 2017. Published at ICLR 2017; arXiv preprint posted in 2016.

Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):160035, 2016. doi: 10.1038/sdata.2016.35.

Bob Kemp, Aeilko H Zwinderman, Bert Tuk, Hilbert AC Kamphuisen, and Josefien JL Oberye. Analysis of a sleep-dependent neuronal feedback loop: the slow-wave microcontinuity of the eeg. *IEEE Transactions on Biomedical Engineering*, 47(9):1185–1194, 2000. doi: 10.1109/10.867928.

Hyung-Chul Lee, Yoonsang Park, Soo Bin Yoon, Seong Mi Yang, Dongnyeok Park, and Chul-Woo Jung. Vitaldb, a high-fidelity multi-parameter vital signs database in surgical patients. *Scientific Data*, 9(1):279, 2022. doi: 10.1038/s41597-022-01411-5.

Attila Reiss, Ina Indlekofer, Philip Schmidt, and Kristof Van Laerhoven. Deep ppg: Large-scale heart rate estimation with convolutional neural networks. *Sensors*, 19(14):3079, 2019. doi: 10.3390/s19143079.

Francesco Scardulla, Gloria Cosoli, Susanna Spinsante, Angelica Poli, Grazia Iadarola, Riccardo Pernice, Alessandro Busacca, Salvatore Pasta, Lorenzo Scalise, and Leonardo D'Acquisto. Photoplethysmograhic sensors, potential and limitations: Is it time for regulation? a comprehensive review. *Measurement*, 218:113150, 2023. doi: 10.1016/j.measurement.2023.113150.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, pp. 368–377, 1999. Conference publication in 1999; arXiv version posted in 2000.

Patrick Wagner, Nils Strodthoff, Ralf-Dieter Bousseljot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. Ptb-xl, a large publicly available electrocardiography dataset. *Scientific data*, 7(1):154, 2020. doi: 10.1038/s41597-020-0495-6.