BeliefPPG: Uncertainty-aware Heart Rate Estimation from PPG signals via Belief Propagation
(Supplementary Material)

Valentin Bieri[*1] **Paul Streli**[*1] **Berken Utku Demirel**[1] **Christian Holz**[1]

[1]Department of Computer Science, ETH Zürich, Switzerland

A TRANSITION FUNCTION SELECTION

While we experimented with alternative distributions for the transition function, including LaPlacian, Gaussian, and Levy distributions of the absolute difference and relative difference in HR between two beat-to-beat intervals, we observed that our transition function using a discretized Gaussian prior fit on the logarithmic change $\log \frac{y_t}{y_{t-1}}$ led to the best results while offering a reasonable fit to the observed histogram of heart range changes (see Figure 1).

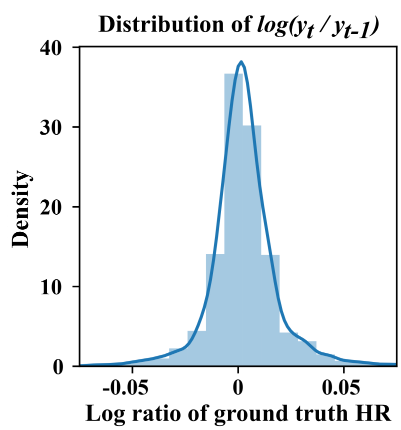

Figure 1: Logarithmic change of HR values in the IEEE dataset.

B SELECTION OF σ_y

We represent the ground-truth HR y_t as a discretized normal distribution with density $\mathcal{N}\left(y_t, \sigma_y^2\right)$. The selection of the hyperparameter σ_y^2 is based on an ablation study, and the results, as demonstrated in Table 1, indicate minimal sensitivity to variations in y_t.

[*]These authors contributed equally to this work.

Table 1: Ablations of BeliefPPG with different σ_y under LoSo-CV on BAMI-1. We report MAE, its standard deviation across subjects (STD-MAE) and the negative log-likelihood (NLL) on the test set.

σ_y	0.25	0.5	0.75	1.0	1.5	2.0	2.5	3.0	4.0
MAE	2.15	2.30	2.34	2.18	**2.00**	2.45	2.29	2.45	2.35
STD MAE	0.90	1.30	1.70	1.00	**1.00**	1.80	1.50	1.50	1.40
NLL	4.13	4.2	4.22	4.13	**4.12**	4.28	4.23	4.3	4.33

C MAPE ACROSS ACTIVITIES

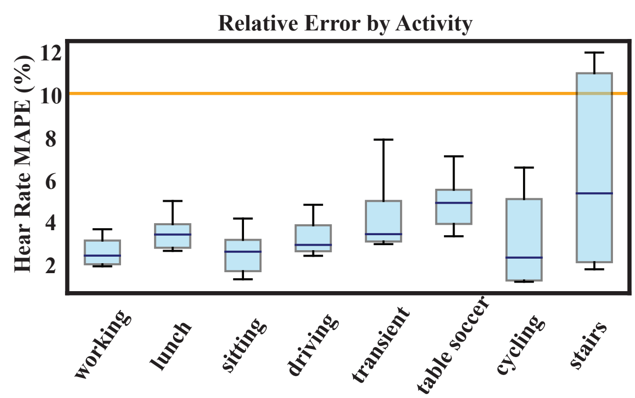

Figure 2: Mean absolute percentage errors (MAPEs) visualized by activity for LoSo on DaLiA. Boxes correspond to interquartile ranges and whiskers to 10th and 90th percentiles over subjects. The blue lines indicate the median, and the orange line marks the acceptable error of 10%.

Figure 2 presents the Mean Absolute Percentage Errors (MAPE) across activities compared to the AAMI standard [ANSI/AAMI, 2002] for the DaLiA dataset. The AAMI standard sets the acceptable limits for HR monitoring within ±10%, which is implemented using the MAPE statistic as defined by the Consumer Technology Association [Association, 2018].

The results demonstrate that the median MAPE for all activities is significantly below this limit. Notably, none of the 15 peak detectors benchmarked by Charlton et al. [Charlton et al., 2022] achieved this level of accuracy.

D ARCHITECTURE DETAILS

Table 2 provides architecture and training details. We implemented our network using TensorFlow 2.8.0.

References

ANSI/AAMI. Cardiac monitors, heart rate meters, and alarms. *Arlington*, 2002.

Consumer Technology Association. Physical activity monitoring for heart rate, ansi/cta 2065, 2018. Accessed: 2023-02-17.

Peter H Charlton, Kevin Kotzen, Elisa Mejía-Mejía, Philip J Aston, Karthik Budidha, Jonathan Mant, Callum Pettit, Joachim A Behar, and Panicos A Kyriacou. Detecting beats in the photoplethysmogram: benchmarking open-source algorithms. *Physiological Measurement*, 43(8): 085007, aug 2022. doi: 10.1088/1361-6579/ac826d.

Table 2: Details about BeliefPPG's heart rate network architecture and training configuration

Heart rate network architecture								
	Up-/Downspl. fac.	Comment	Kernel Size	Padding	Filters / Inner Dim	Activation	Dropout	Output shape
Time-Frequency Branch *// input shape:* $(W_s, N_s, 2) = (7, 64, 2)$								
2x Conv2D	-	-	3x3	same	32	leaky_relu	0.1	(7,64,32)
embedding 4x	-	-	-	-	32	-	-	4 x (7,64,32)
2x attention + reduce mean	-	reduces 1^{st} axis	-	-	-	-	-	(64, 32)
3x downsampling block (1D)	4	stride=poolsize=4	3x1	same	12, 24, 48	relu	0.2	(1, 48)
bottleneck attention	-	-	2x1	-	48	tanh, relu	0.2	(1, 48)
3x upsampling block	4	upspl. size=4	3x1	same	48, 24, 12	relu	0.2	(64, 12)
Conv1D	-	-	1x1	same	1	softmax	-	(64,)
Time Branch *// input shape:* $(L_x, 1) = (1280, 1)$								
Conv1D + bn + MaxPool	4	dilation_rate=2	10	causal	16	leaky_relu	0.1	(320, 16)
Conv1D + bn + MaxPool	4	dilation_rate=2	10	causal	16	leaky_relu	0.1	(80, 16)
LSTM	-	-	-	-	64	tanh	0.1	(80, 64)
LSTM	-	-	-	-	64	tanh	0.1	(64,)
Dense 2x	-	-	-	-	48	leaky_relu	-	2x (1, 48)
Training Parameters								
Batch Size	128							
Optimizer	Adam(lr=0.00025)							
LR Scheduler	ReduceLROnPlateau(factor=0.5, min_lr=1e-10 monitor="loss", patience=3)							
Stopping criterion	EarlyStopping(patience=40 restore_best_weights=True, monitor="val_loss")							
TTA: gaussian noise std	0.25							
TTA: max stretching factor	25%							
σ_y: label standard deviation	1.5							