

# KID-PPG: Knowledge Informed Deep Learning for Extracting Heart Rate from a Smartwatch

Christodoulos Kechris, Jonathan Dan, Jose Miranda, and David Atienza, *Fellow, IEEE*

**Abstract**—Accurate extraction of heart rate from photoplethysmography (PPG) signals remains challenging due to motion artifacts and signal degradation. Although deep learning methods trained as a data-driven inference problem offer promising solutions, they often underutilize existing knowledge from the medical and signal processing community. In this paper, we address three shortcomings of deep learning models: motion artifact removal, degradation assessment, and physiologically plausible analysis of the PPG signal. We propose KID-PPG, a knowledge-informed deep learning model that integrates expert knowledge through adaptive linear filtering, deep probabilistic inference, and data augmentation. We evaluate KID-PPG on the PPGDalia dataset, achieving an average mean absolute error of 2.85 beats per minute, surpassing existing reproducible methods. Our results demonstrate a significant performance improvement in heart rate tracking through the incorporation of prior knowledge into deep learning models. This approach shows promise in enhancing various biomedical applications by incorporating existing expert knowledge in deep learning models.

**Index Terms**—Photoplethysmography, Heart Rate, Motion Artifacts, Acceleration, Source Separation, Deep Learning, Knowledge Informed AI

## I. INTRODUCTION

Photoplethysmography (PPG) is a non-invasive technique used to optically acquire the Blood Volume Pulse (BVP) [1]. Its widespread adoption and ease of integration into wearable devices, especially smartwatches, have made PPG a popular choice for continuous and unobtrusive heart rate monitoring compared to electrocardiography (ECG). However, movement can introduce significant artifacts into PPG signals, complicating signal interpretation. These motion artifacts (MA) can overlap the actual BVP signal, further complicating their removal [2]. To address this challenge, numerous methods have been proposed to estimate the heart rate (HR) of PPG corrupted by MA. These methods generally fall into two categories: Signal Processing (SP) and Deep Learning (DL) approaches.

SP methods focus mainly on isolating the BVP component and minimizing the impact of MA [2]. These methods often use a motion reference signal acquired from sensors such as accelerometers or gyroscopes to aid in MA removal [3]–[7]. Acceleration and angular velocity are generally agreed to

be effective references for periodic motion but offer limited correlation with PPG signals during random movements [3]. Once the motion components are filtered out, HR is typically extracted from the filtered PPG signal by identifying its principal frequency component [3]–[11], which is then attributed to the heart activity.

DL offers an alternative approach by combining filtering and HR estimation within a single model. Similarly to SP methods, DL methods have used acceleration as a reference signal for motion. In these models, both PPG and acceleration are given directly as inputs to the network and are fused by the model to produce a point estimate of HR [12]–[14] or an HR distribution [15], [16]. The network is typically trained in a supervised manner using synchronized ECG-derived HR as ground truth labels. Several advanced DL techniques, including attention [13] and data augmentation [17], have been used for PPG-based HR extraction.

Existing DL methods typically approach PPG-based HR estimation as a purely data-driven inference task, overlooking valuable prior knowledge from the medical and signal processing fields regarding BVP and PPG. Integrating task-specific prior knowledge into machine learning models has emerged as a promising strategy to improve explainability, robustness, and generalizability, particularly in scenarios with limited available data [18].

In this work, we explore the integration of prior knowledge into DL models for PPG-based HR inference. We identify failure cases of current DL models and propose three mechanisms to integrate prior knowledge into DL models effectively addressing these shortcomings. The resulting approach, which we term KID-PPG, represents a knowledge-informed DL-based HR inference model. To conduct our analysis, we take advantage of the publicly available PPGDalia dataset [12]. Our code for the experiments is available here: <https://github.com/esl-epfl/KID-PPG-Paper>. Through this investigation, we shed new light on the efficacy of DL models in processing PPG signals affected by MA and provide valuable insights into the recoverability of the BVP under challenging MA conditions. Our contributions include:

- We identify three key factors contributing to erroneous HR estimations: DL models fail to separate MA from BVP, infer out-of-distribution HR samples, and estimate HR in sample affected by catastrophic MA.
- We address these limitations by incorporating prior knowledge into the DL workflow: explicitly defined MA separation task, guided probabilistic inference, and data

augmentation.

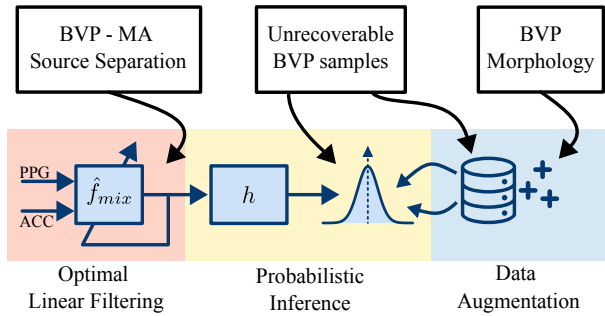
- We design KID-PPG, a DL model for HR inference, achieving a mean absolute error (MAE) of 2.96 beats per minute (BPM) on PPGDalia. KID-PPG also provides estimates of uncertainty as a proxy for the assessment of the severity of the BVP artifact along with HR inference.
- We provide an open-source package to estimate HR from PPG and acceleration signals to the research community, which allows to repeat all our experiments and move forward the field of PPG analysis on embedded devices and wearables. The package is available here: <https://github.com/esl-epfl/KID-PPG>.

## II. METHODOLOGY

We have identified the following three key points underutilized in existing DL models for HR extraction:

- 1) Robust HR tracking requires MA removal.
- 2) In some cases, BVP can be degraded by MA to the extent that it is unrecoverable.
- 3) The BVP has specific morphology and characteristics.

These are incorporated into our DL models through three mechanisms: Explicit Source Separation, Guided Probabilistic Inference and Data Augmentation. An overview of our methodology is presented in Fig. 1.



**Fig. 1:** Knowledge-Informed deep-learning (DL) for heart-rate extraction (KID-PPG) incorporates prior knowledge on motion artifacts (MA), unrecoverable blood volume pulse (BVP) samples, and BVP morphology into DL models through three mechanisms: linear filtering, probabilistic inference, and data augmentation. The input of KID-PPG consists of a PPG signal along with an accelerometer signal. The linear filter ( $\hat{f}_{mix}$ ) separates the BVP from MA to produce a filtered PPG signal, which serves as an input to the DL model ( $h$ ). The model uses probabilistic inference to assess the degradation of the PPG signal and uses data augmentation to better characterize the BVP morphology.

### A. Motion Artifact Removal

An inherent challenge in HR inference based on PPG is the mixing of MA components ( $x_{MA}(t)$ ) with heart-related BVP ( $x_{bvp}(t)$ ) through a mixing process denoted as  $f_{mix}$ :

$$x_{ppg}(t) = f_{mix}(x_{bvp}(t), x_{MA}(t), t) \quad (1)$$

As is common practice, HR is inferred on 8-second windows with a 2-second overlap [4], [5], [12]–[16], [19]–[21]. Hence,

the  $i$ -th 8-second PPG sample is denoted as the  $(N \times 1)$  vector  $\mathbf{x}_{ppg_i} = [x_{ppg}(t_i), x_{ppg}(t_i + \Delta t), \dots, x_{ppg}(t_i + 8)]$ , where  $\Delta t = \frac{1}{f_s}$ ,  $f_s$  is the sensor's sampling frequency, and  $N = 8sec \cdot f_s$  samples. Similarly, the corresponding BVP is denoted as  $\mathbf{x}_{bvp_i}$ . The 3-axis ACC samples form the  $(N \times 3)$  matrix  $X_{acc_i} = [\mathbf{x}_{acc_x_i}, \mathbf{x}_{acc_y_i}, \mathbf{x}_{acc_z_i}]$ . Robust HR inference depends on the heartbeat-related  $\mathbf{x}_{bvp}$  [4]. However, MA can significantly distort the morphology of BVP in the PPG observation (Eq. 1) [5]. Many SP methods employ source separation to approximate the unmixing process  $f_{mix}^{-1}$ , allowing the HR inference module to rely on an approximation  $\hat{x}_{bvp}(t) \approx x_{bvp}(t)$  [3]–[5], [19], thus mitigating the effect of MA.

In contrast, training a deep end-to-end HR estimator,  $g(\cdot)$ , with inputs  $[\mathbf{x}_{ppg_i}, X_{acc_i}]$ , for the HR inference task, [12]–[16], [20], [21], does not guarantee modeling of the unmixing process, namely:

$$g = \hat{f}_{mix}^{-1} \circ h \quad (2)$$

where  $h(\cdot)$  is an estimator of HR after component unmixing. Instead, the network may employ various fusion strategies  $g = f_{fusion} \circ h'$ , where  $h'$  is an HR estimator after the two modalities have been fused by  $f_{fusion}$ . The convergence criterion for  $f_{fusion}$  is the minimization of the HR loss, e.g. the MAE. Thus, this approach does not ensure BVP-based HR inference and may allow the network to learn spurious relations with the motion signals [22].

Therefore, we propose the inclusion of an explicitly defined source separation task to disentangle the BVP component from MA. Without loss of generality, we model the motion artifact mixing  $f_{mix}$  as a linear process [7], [23], although other approaches are available [24], [25]. Furthermore, we assume that  $f_{mix}$  is a stationary process and that hand acceleration is a suitable reference signal for  $x_{MA}(t)$  [3]. Hence, Eq. 1 can be written as:

$$x_{ppg}(t) = x_{bvp}(t) + A_{mix} * [x_{acc_x}(t), x_{acc_y}(t), x_{acc_z}(t)] + noise(t) \quad (3)$$


where  $A_{mix}$  is a spatio-temporal mixing filter,  $x_{acc_x}$ ,  $x_{acc_y}$ ,  $x_{acc_z}$  are the accelerations of the hand on the corresponding axis, and  $*$  is the convolution operation. Separating the motion artifacts involves estimating the approximation  $\hat{A}_{mix} \approx A_{mix}$ .

To estimate  $\hat{A}_{mix}$  we use a linear two-layer convolutional network denoted as  $\hat{f}_{mix}$  (Fig. 2). The first convolutional layer applies linear spatio-temporal filtering on the 3-channel accelerometer signals, while the second layer merges the three channels into one MA estimation.

Similarly to [7], we train  $\hat{f}_{mix}$  in an unsupervised manner on pairs  $(X_{acc_i}, \mathbf{x}_{ppg_i})$  using an adaptive filter loss function, denoted as  $\mathcal{L}_{adapt}$ :

$$\mathcal{L}_{adapt} = MSE \left( FFT\{\hat{f}_{mix}(X_{acc_i})\}, FFT\{\mathbf{x}_{ppg_i}\} \right) \quad (4)$$

Here,  $MSE$  represents the Mean Squared Error, defined as  $\mathbb{E}[Error^2]$ , and  $FFT\{\cdot\}$  denotes the Fast Fourier Transform.



**Fig. 2:** Linear model for separating the motion artifacts (MA) from the blood volume pulse (BVP) components in the PPG signal. A linear two-layer convolutional network takes the three axis of the accelerometer signal as an input ( $\mathbf{x}_{acc}$  along with the PPG ( $\mathbf{x}_{ppg}$ ) to produce a filtered PPG signal ( $\hat{\mathbf{x}}_{bvp}$ ).

Once the training converges, we use the resulting  $\hat{f}_{mix}$  to estimate  $\hat{\mathbf{x}}_{bvp} \approx \mathbf{x}_{bvp}$ :

$$\hat{\mathbf{x}}_{bvp} = \mathbf{x}_{ppg} - \hat{f}_{mix}(X_{acc}) \quad (5)$$

The estimated  $\hat{\mathbf{x}}_{bvp}$  is then inputted into a DL model, denoted as  $h$  (Eq. 2).

### B. Temporal Attention Model

A convolutional neural network based on [13] is used to extract an embedding  $W_i$  for each sample  $\hat{\mathbf{x}}_{bvp_i}$ .

HR temporal relationships are usually modeled as a smoothing filter during post-processing, reducing temporal granularity, or with LSTMs, increasing the model's complexity. We propose to model the temporal relationship between two consecutive samples as an attention operation. This approach allows us to consider the progression of the PPG signal embeddings while maintaining temporal granularity and computational simplicity. Let  $E_{i-1}, E_i$  denote the embeddings of two consecutive PPG frames. The multi-head temporal attention operation can be defined as [13], [26]:

$$Attention_{Temp}(E_i, E_{i-1}) = softmax\left(\frac{E_i E_{i-1}^T}{\sqrt{d}}\right) E_{i-1} \quad (6)$$

where  $d$  is the dimensionality of the embedding. Additionally, we incorporate a residual connection [26], [27]:  $E_i + Attention_{Temp}(E_i, E_{i-1})$ .

### C. Guided Probabilistic Heart Rate Extraction

The DL models discussed in [12]–[14] and Sub-sections II-A and II-B produce a point estimate of HR, implying that each sensor readout value contains a BVP component which can be isolated. However, in real-world conditions, the BVP component might be degraded beyond the point of reconstruction, leaving the MA component as the sole source of information in the PPG sample, Eq. 1. If  $x_{bvp}(t)$  is not observed, the model performs HR extraction on irrelevant information, Eq. 1. This is illustrated in Fig. 3 where the BVP component is severely degraded, yet a point-estimate DL model continues to infer HR (Q-PPG [14]).

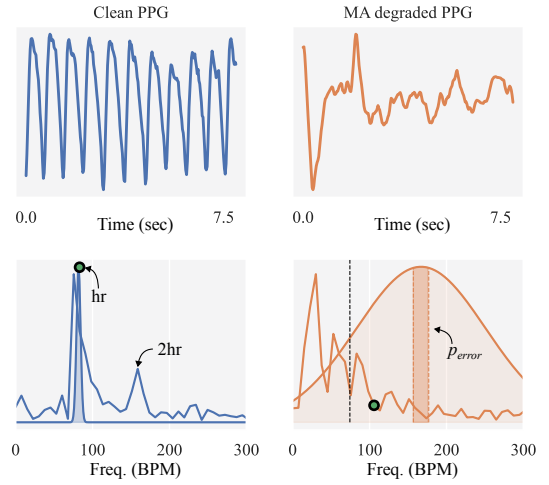
To address this issue, we propose to design the model as a probability estimator of  $HR$ . Specifically, we choose the

normal distribution parameterized as  $HR \sim \mathcal{N}(\mu_{hr}, \sigma_{hr}^2)$ . Inspired by [16] and [28], we model the heteroscedastic aleatoric uncertainty using a two-unit fully connected output to represent both  $\mu_{hr}$  and  $\sigma_{hr}$ . If no heart-related information is available in the PPG sample, then any physiologically valid heart rate is possible, therefore, the HR estimator should produce a large uncertainty, as depicted in Fig. 3.

The estimated HR distribution can also be used as an error classifier. We define the error classification probability as

$$\begin{aligned} p_{error}(Thr | \mathbf{x}_{ppg}, \mu_{hr}, \sigma_{hr}) \\ = P(\mu_{hr} - Thr < HR < \mu_{hr} + Thr | \mu_{hr}, \sigma_{hr}) \\ = F_{\mu_{hr}, \sigma_{hr}}(\mu_{hr} + Thr) - F_{\mu_{hr}, \sigma_{hr}}(\mu_{hr} - Thr) \end{aligned} \quad (7)$$

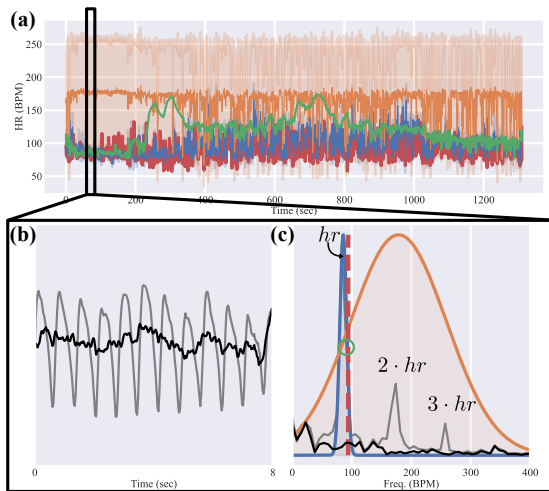
where  $F_{\mu_{hr}, \sigma_{hr}}(x)$  is the cumulative distribution of the Gaussian distribution  $\mathcal{N}(\mu_{hr}, \sigma_{hr}^2)$ :  $F_{\mu_{hr}, \sigma_{hr}}(x) = \Phi\left(\frac{x - \mu_{hr}}{\sigma_{hr}}\right)$ . This error classifier predicts a high probability of the HR estimation being untrustworthy (error  $\geq Thr$ ) when  $p_{error} \leq 0.5$ .



**Fig. 3:** Probabilistic HR inference example on a clean (left) vs MA degraded PPG sample (right). The top row shows the raw PPG data. The bottom row shows the FFT of the PPG. The example is taken from S6 of PPGDalia. The green circles indicate the true ECG HR for the two samples. The probability density functions of  $\mathcal{N}(\mu_{hr}, \sigma_{hr}^2)$  are overlaid on the frequency representations of the corresponding PPG sample. The error classification probability,  $p_{error}$  is also illustrated. In the right sample the non-probabilistic point estimate HR inference of DL model is represented as a black vertical line (Q-PPG).

To improve the robustness of our error classifier, we guide the network to base its HR inference on the BVP component, as relying on HR-inference loss only can lead to spurious behavior. Fig. 4 illustrates this through an example taken from PPGDalia (S6), where both the probabilistic and the Q-PPG models continue to infer HR even after filtering the BVP component out of the original PPG signal (MAE 4.11 BPM and 3.33 BPM accordingly). The probabilistic model is overconfident, with the  $p_{error}(Thr = 10BPM)$  classifier dropping only 7.44% of the samples. We tackle this limitation by generating realistic samples in which the BVP component

is extremely degraded. It has been empirically observed that in PPG recordings, BVP information is mainly located at the HR frequency, and its second and third harmonics [19], [29]. Therefore, for each sample  $(\mathbf{x}_{ppg_i}, hr_i)$  in the original training dataset, we filter  $\mathbf{x}_{ppg_i}$  by band-stopping (Finite Input Response with 81 taps) the frequencies around  $hr_i$ ,  $2hr_i$ ,  $3hr_i$  (cutoff from  $i \cdot hr - 2.5BPM$  to  $i \cdot hr + 2.5BPM$ ), forming a new signal  $\mathbf{x}_{noise_i}$  (Fig. 4 b, c). Using the original  $\mathbf{x}_{ppg_i}$  as the seed for  $\mathbf{x}_{noise_i}$  helps maintain realism. Since we need a ground truth label for the supervised training task, we create a random HR label:  $hr_{noise_i} \sim Uniform(40, 300)$ . The lowest HR value is set to account for slower HR [30] and the highest reflects the theoretical maximum human HR [31]. The pair  $(\mathbf{x}_{noise_i}, hr_{noise_i})$  is then added to the set of auxiliary adversarial examples. 50% of the adversarial examples set is randomly sampled to form the *adversarial-examples-subset*.



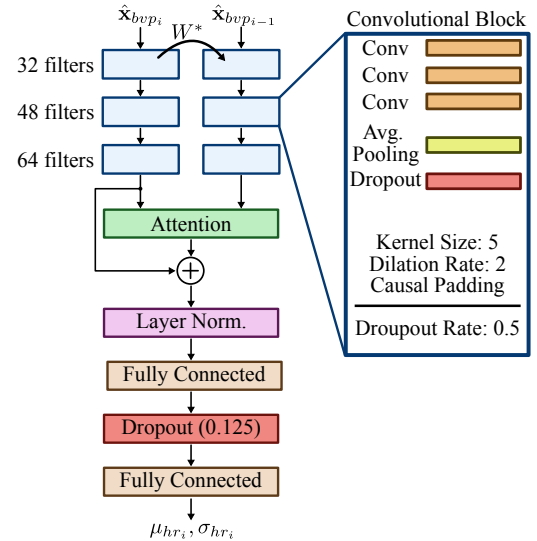
**Fig. 4:** Inference after filtering the BVP component out of the PPG. Example taken from S6 of PPGDalia. (a) Inferences across the entire session from Q-PPG (red), probabilistic (blue) and guided probabilistic (orange). For the probabilistic models, the range of one standard deviation is also presented. (b) Sample example with initially clean PPG (grey) and synthetically degraded (black). (c) Frequency domain representation of the example sample and HR inferences from the three models. True HR is presented with a green circle.

In summary, by adopting a probabilistic approach to HR estimation and guiding the network to focus on the BVP component, we aim to improve the reliability of HR inference from PPG signals. The full KID-PPG network configuration is presented Fig. 5.

#### D. Data Augmentation

We propose a data augmentation scheme to address the limited number of high HR samples in the available datasets, similarly to [17]. We generate synthetic PPG waveforms corresponding to higher heart rate frequencies, forming the *high-heart-rate* training subset. The following procedure is followed:

- 1) Locate the 8-sec. samples designated as clean PPG for



**Fig. 5:** KID-PPG network architecture.  $W^*$  indicates weight sharing between the convolution blocks for the  $\hat{\mathbf{x}}_{bvp_i}$  and  $\hat{\mathbf{x}}_{bvp_{i-1}}$  branches.

which the main frequency component is close to the ground truth HR.

- 2) Artificially speed up the sample by x2.
- 3) Discard samples with  $HR \geq 300BPM$  ([31]).

The original training dataset, the *high-heart-rate* and *adversarial-examples-subset* subsets are then merged into the final training set.


### III. EXPERIMENTAL SETUP

For our experiments we use the publicly available PPGDalia data [12]. This dataset comprises synchronized ECG and wristworn PPG and acceleration recordings from 15 subjects, with approximately two-hour recording sessions per subject. A systematic temporal shift between PPG and ACC was identified and manually corrected. The ECG-based HR provided in the dataset serves as the ground truth. We adopt the leave-one-subject-out cross-validation procedure proposed in [12] for all experiments.

For the MA removal step, we train an adaptive model separately for each subject and activity given that converging  $\hat{f}_{mix}$  using  $\mathcal{L}_{adapt}$ , Eq. 4, requires stationarity of the  $f_{mix}$  process. We employ the MAE loss function and the stochastic gradient descent (SGD) to converge the model ( $lr = 1e - 7$ ,  $momentum = 1e - 2$ ).

To evaluate the MA removal step, we employ the Q-PPG [14], Attention Model [13] and Temporal Attention model (Sub-section II-B) as HR inference models,  $h$  Eq. 2. For the Q-PPG and Attention models we used the hyperparameters reported in [14] and [13], respectively. The Temporal Attention model was trained using Adam optimizer ( $lr = 0.0005$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1e - 08$ ).

The probabilistic models were trained using the negative log-likelihood (NLL) loss and the same training strategy used for the non-probabilistic ones. NLL is also used to evaluate their performance. Furthermore, we use the True Positive Rate



**Fig. 6:** Estimating HR on stairs activity of S6 with Q-PPG (top row - blue) and with Adaptive filtering and Q-PPG (bottom - orange). The original Q-PPG (blue) takes the PPG and accelerometer (ACC) as input. Adaptive + Q-PPG (orange) takes the output ( $b_{vp}$ ) of the adaptive filter ( $f_{mix}$ ) as input.

(TPR) and the F1 score to evaluate the classifier  $p_{error}$  for its ability to correctly identify untrustworthy samples. We consider a positive classification when the error classifier predicts a high probability of error,  $p_{error}(Thr) \leq 0.5$ .

A summary of all experiments is presented in Table I.

<b>Base models</b>	
• Q-PPG	Introduced in A. Burrello et al. [14].
• Attention	Introduced in P. Kasnesis et al. [13].
<b>Adaptive</b>	
• Adaptive + Q-PPG	Sub-section II-A
• Adaptive + Attention	
<b>Data augmentation</b>	
• Adaptive + Attention + HR Aug.	Sub-section II-D
<b>Probabilistic</b>	
• Probabilistic model	Adaptive + Attention model adapted for probabilistic inference (Sub-section II-C).
• Probabilistic Temporal Attention model	Probabilistic model with temporal attention (Sub-section II-B).
• KID-PPG	Probabilistic model with temporal attention, guided probabilistic training and High-HR augmentation (Sub-section II-C).


**TABLE I:** Summary of evaluated models.

## IV. RESULTS

### A. Results on Motion Artifact Removal

The adaptive linear filter effectively removes MA that are linearly coupled with acceleration signals. An illustrative example is provided in Fig. 6 for subjects S6 during an episode where the subject is walking up and down stairs. The DL model alone fails to decouple motion information, resulting in significant errors (activity MAE 13.66 BPM). In contrast, the Adaptive + DL Model combination substantially reduces the error (activity MAE 2.74 BPM). In particular, during periods when Q-PPG loses track of the HR, the BVP signal, located at  $hr$  and  $2 \cdot hr$ , remains discernible. The loss of HR tracking is not attributed to severe degradation of BVP due to MA.

Similarly, for S5 (Fig. 7), the BVP component is visually identifiable, yet the DL model fails to disentangle it from MA



**Fig. 7:** Effect of the MA Removal for the stairs activity of subjects S5 (top), and S9 (bottom) with Q-PPG (blue) and Adaptive + Q-PPG (orange). Ground truth HR is represented in green.

(activity MAE 69.96 BPM). Our filtering approach effectively isolates the BVP, resulting in lower HR inference error (activity MAE 24.96 BPM). However, during certain periods, 20-40sec. and 80-100sec., the DL model loses track of the HR. We discuss this further in Sub-section IV-B.


For S9 (Fig. 7) the situation differs slightly. MA removal still outperforms the DL model, with Q-PPG stairs activity MAE 16.91 BPM vs. MAE 10.37 BPM for the Adaptive + Q-PPG. However, in this case, MA degradation is more severe, and the MA component is not linearly coupled with ACC.

Consequently, the DL model sometimes loses track of the HR even after the MA removal step. This limitation is addressed by the error classifier.

The overall per-subject results are presented in Table II. Both Q-PPG and the Attention models perform better when data are preprocessed by the adaptive step, with MAE 4.81 BPM (Q-PPG) vs. 3.98 (Q-PPG with adaptive step) and MAE 4.44 vs. 3.70 for the Attention model. Notably, adding a preprocessing step to Q-PPG outperforms the original Attention model. For all but one subject, S9, the best performing model is obtained by employing the adaptive preprocessing step (Table II).

### B. High-HR Augmentation

Training with High-HR augmentation results in a higher MAE (3.79 BPM) compared to not augmenting (3.70 BPM). Subjects S1, S5, S8 and 13 benefited from the augmentation (Table II). High-HR augmentation teaches the model to infer HR in a broader range than that of the original dataset (Fig. 8). In Fig. 8 the HR inference error for S5 during Stairs activity decreases MAE from 25.33BPM to 6.73BPM. Consequently, for the augmented model, there is a wider range of MA components that can be mistakenly inferred as HR. Additionally, the expected HR is now artificially inflated, leading to higher errors in considerably degraded samples in which HR inference reflects the HR distribution encoded in the model weights. As a result, the model's overall MAE is increased. This limitation is addressed by the probabilistic mechanism (Sub-section IV-C).

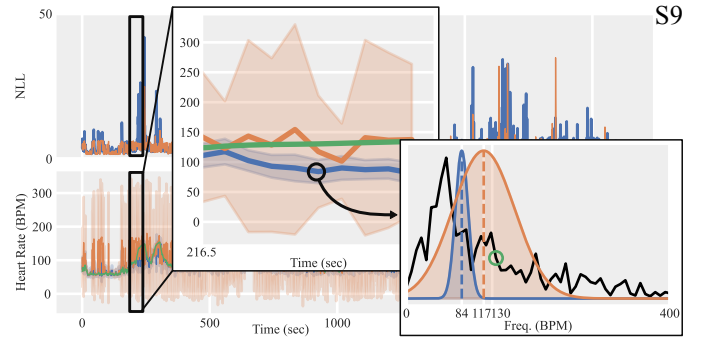


**Fig. 8:** Effect of out-of-distribution samples on the model inference. (a) HR inference of the model with (orange) and without (blue) High-HR augmentation. The maximum HR value in the training set is presented with a red dashed line. (b) Example sample in the frequency domain along with model prediction and ground truth HR (green circle). The HR distributions of the non-augmented and High-HR augmented training sets are also analyzed in the corresponding colors.

### C. Guided Probabilistic Inference


The NLL losses for the probabilistic models are summarized in Table III. Our proposed methodology, KID-PPG, significantly outperforms BeliefPPG [15] (NLL 4.78) and reduces NLL by approximately 40%.

Although the non-augmented and augmented probabilistic temporal attention models achieve similar mean NLL, they



**Fig. 9:** Example of probabilistic output of augmented (orange) vs non-augmented model (blue) for subject S9 and ground truth HR (green circle). For the example, its frequency domain representation verifies that the PPG signal is considerably corrupted.

differ in their ability to infer untrustworthy samples. The guided probabilistic model dropped 98.58% of the samples in the example in Fig. 4. A real data example is presented in Fig. 9. The non-augmented NLL presents "spikes" or outliers corresponding to overconfident but wrong HR estimations, with the highlighted sample showing a 46.30 BPM Absolute Error with 9.72 BPM STD. Adversarial augmentation acts as a probabilistic regularization, reducing the networks' overconfidence (12.98 BPM Absolute Error with 46.92 BPM STD).



**Fig. 10:** TPR and F1-score for KID-PPG (Guided Probabilistic training and High-HR augmentation), orange, vs Unguided probabilistic model (Unguided probabilistic training and no High-HR augmentation), blue.

The overall performance of the error classifiers for KID-PPG and the unguided probabilistic models, TPR and F1-Score, is depicted in Fig. 10. Selecting a high threshold ( $Thr = 10BPM$ ), KID-PPG achieves an average MAE of 2.85 BPM.

## V. DISCUSSION

Our analysis provides novel insights into the challenges of PPG-based HR inference, particularly within the PPG-Dalia dataset. The difficulties encountered by DL model to accurately estimate HR have often been attributed to out-of-distribution high HR samples [12], [21]. Specifically, S5 has been highlighted as a particularly challenging subject [12], [17], [21]. Our findings emphasize three key factors

	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	Avg
<b>Signal Processing</b>																
SpaMaPlus [12]	8.86	9.67	6.4	14.1	24.06	11.34	6.31	11.25	16.04	6.17	15.15	12.03	8.50	7.76	8.29	11.06
TAPIR [10]	4.50	4.50	3.20	6.00	5.00	3.40	2.80	6.30	8.00	2.90	5.10	4.70	3.10	5.00	4.10	4.57
CurToSS [11]	5.40	4.30	3.00	8.00	<b>2.20</b>	2.80	3.30	8.50	12.60	3.60	3.60	6.10	3.00	5.50	3.70	5.04
<b>Deep Learning</b>																
DeepPPG [12]	7.73	6.74	4.03	5.90	18.51	12.88	3.91	10.87	8.79	4.03	9.22	9.35	4.29	4.37	4.17	7.65
NAS-PPG [20]	5.46	5.01	3.74	6.48	12.68	10.52	3.31	8.07	7.91	3.29	7.05	6.76	3.84	4.85	3.57	6.16
Q-PPG [14]	4.29	3.62	2.44	5.73	10.33	5.26	2.00	7.09	8.6	3.09	4.99	6.25	1.92	3.02	3.55	4.81
TEMPONet	4.37	3.74	2.43	5.49	13.48	5.71	2.23	7.86	8.94	3.32	5.34	7.71	2.03	2.94	3.58	5.27
Augmentation [17]	4.97	4.34	2.39	6.14	9.41	3.63	2.23	9.14	10.98	3.4	5.27	7.64	2.05	2.84	3.61	5.20
TimePPG	4.51	3.37	2.33	5.25	14.68	4.76	2.37	8.04	8.75	3.30	5.19	8.08	2.29	3.02	3.49	5.29
Q-PPG [14]	4.29	3.62	2.44	5.73	10.33	5.26	2.00	7.09	8.60	3.09	4.99	6.25	1.92	3.02	3.55	4.81
Attention [13]	4.75	3.31	2.22	5.25	7.43	4.22	2.28	8.93	6.95	2.93	3.98	6.57	1.70	3.22	2.88	4.44
<b>Ours</b>																
Adaptive + Q-PPG	3.80	3.50	2.07	5.18	5.76	3.38	1.59	6.99	8.94	2.75	3.37	5.61	1.49	2.45	2.80	3.98
Adaptive + Attention	4.55	3.12	1.88	5.27	4.00	<b>2.54</b>	<b>1.28</b>	7.20	8.60	2.61	3.06	4.91	1.50	<b>2.28</b>	2.64	3.70
Adaptive + Attention + HR Aug	4.27	3.46	2.07	5.61	3.01	2.74	1.39	7.13	9.53	2.77	3.58	4.52	<b>1.48</b>	2.48	2.84	3.79
<b>Ours - Probabilistic</b>																
KID-PPG ( $Thr = 10BPM$ )	<b>2.99</b>	<b>2.54</b>	<b>1.84</b>	<b>3.83</b>	2.8	2.56	1.50	<b>5.16</b>	<b>4.67</b>	<b>2.33</b>	<b>2.90</b>	<b>3.73</b>	1.72	2.33	<b>1.91</b>	<b>2.85</b>

TABLE II: Mean Absolute Error performance of point-estimate HR and probabilistic (after retention) models on the PPGDalia dataset. All values are in Beats-per-Minute.

Method	Mean Negative Log-Likelihood
BeliefPPG [15]	4.78
Adapt + Attention + Prob	3.25
Adapt + Temp. Attention + Prob	3.08
KID-PPG	<b>2.99</b>

TABLE III: Summarized NLL results for the proposed probabilistic models.

contributing to erroneous HR estimations: DL models fail to separate MA from BVP, out-of-distribution HR samples and catastrophic MA.

**DL models encounter difficulties in separating linear MA.** The introduction of an explicit linear MA separation filter notably improved model performance, enabling HR inference in cases where the DL model alone had failed. This finding suggests a possible physics-based explanation, as hinted at by Lee et al. [3], who empirically linked wrist-worn PPG motion artifact components to the physical forces acting on the hand during physical activities.

**Out-of-distribution HR samples** have posed challenges, particularly exemplified by subject S5, frequently cited in the literature for its higher HR samples. Previous studies have reported improvement through data augmentation, particularly in S5 and S6 [17]. In our experiments, mainly subjects S1 and S5 benefited from increased HR. A comparison between S5 and S6 reveals differing underlying causes of high errors, with S6 primarily affected by the PPG - MA coupling, and S5 impacted by both MA and out-of-distribution samples.

**Catastrophic MA** remains a significant challenge, particularly evident in subjects S8 and S9, which consistently exhibit high MAE. Despite various proposed solutions, no one has achieved a low MAE for these two subjects. Furthermore, the rejection of a significant portion of their samples by KID-PPG (27.48% and 53.16% of their samples respectively ( $Thr = 10BPM$ )) indicates severe BVP corruption. This underscores the need for a probabilistic HR estimator and a robust error classifier based on the morphology of the BVP components.

While our work has demonstrated the efficacy of explicitly defining the source-separation task, real-world deployment scenarios, which further explore the assumption of stationarity of  $f_{mix}$ , Eq. 5 are required. Deploying this solution requires identifying when  $f_{mix}$  needs to be updated, potentially through an activity/context recognition module. In addition, other implementations of the source-separation task warrant investigation.

In integrating BVP-morphology-related prior knowledge into KID-PPG, we have proposed a method for creating realistic PPG signals with extreme BVP degradation. Exploring alternative approaches to generate synthetic MA-affected PPG samples and expanding the understanding of MA dynamics with the BVP could inform the design of more realistic augmentation techniques. Evaluating MA degradation presents an additional challenge, as HR-inference accuracy may not fully reflect the level of degradation. Addressing these challenges will be crucial to advancing the robustness and reliability of PPG-based HR inference systems.

## VI. CONCLUSION

In this study, we have introduced a novel method, called KID-PPG, to enhance DL models with expert knowledge integration for HR inference from PPG inputs (combined with a motion reference signal). In particular, we have proposed three main knowledge integration mechanisms: Adaptive Linear Filtering, Guided Probabilistic Inference, and Data Augmentation. Adaptive Filtering removes linear MA, enabling the DL model to infer HR accurately even in segments with high MA, thus strongly reducing MAE compared to existing DL models. Probabilistic inference allows the model to assess the BVP degradation in input signals due to artifacts. Using the proposed error classifier, the model can selectively retain samples with a high probability of maintaining BVP information, improving overall inference accuracy. Our guided probabilistic training strategy, tailored to the morphology of BVP, substantially improved the robustness of the error classifier. Data Augmentation extended the range of HR that

can be reliably inferred by the model, further enhancing its performance. This work highlights the important benefits of integrating expert knowledge into DL models. All in all, KID-PPG achieves an overall MAE of 2.85 BPM ( $Thr = 10BPM$ ) on PPGdalia, outperforming all reproducible methods. Our findings demonstrated the efficacy of our approach in advancing HR inference from PPG signals, paving the way for improved healthcare monitoring and diagnostics.

### ACKNOWLEDGMENT

This research was partially supported by the PEDESITE Swiss NSF Sinergia project (GA No. CRSII5\_193813 / 1), the RESoRT project (GA No. REG-19-019) from the Botnar Foundation, and by the Wyss Center for Bio and Neuro Engineering through funding for ESL-EPFL in the Non-invasive Neuromodulation of Subcortical Structures project of the Lighthouse Partnership Agreement with EPFL.

### REFERENCES

- [1] J. Allen, "Photoplethysmography and its application in clinical physiological measurement," *Physiological measurement*, vol. 28, no. 3, p. R1, 2007.
- [2] D. Biswas, N. Simões-Capela, C. Van Hoof, and N. Van Helleputte, "Heart rate estimation from wrist-worn photoplethysmography: A review," *IEEE Sensors Journal*, vol. 19, no. 16, pp. 6560–6570, 2019.
- [3] H. Lee, H. Chung, H. Ko, A. Parisi, A. Busacca, L. Faes, R. Pernice, and J. Lee, "Adaptive scheduling of acceleration and gyroscope for motion artifact cancellation in photoplethysmography," *Computer Methods and Programs in Biomedicine*, vol. 226, p. 107126, 2022.
- [4] Z. Zhang, "Heart rate monitoring from wrist-type photoplethysmographic (ppg) signals during intensive physical exercise," in *2014 IEEE global conference on signal and information processing (GlobalSIP)*. IEEE, 2014, pp. 698–702.
- [5] Z. Zhang, Z. Pi, and B. Liu, "Troika: A general framework for heart rate monitoring using wrist-type photoplethysmographic signals during intensive physical exercise," *IEEE Transactions on biomedical engineering*, vol. 62, no. 2, pp. 522–531, 2014.
- [6] K. Xu, X. Jiang, and W. Chen, "Photoplethysmography motion artifacts removal based on signal-noise interaction modeling utilizing envelope filtering and time-delay neural network," *IEEE Sensors Journal*, vol. 20, no. 7, pp. 3732–3744, 2019.
- [7] D. Yang, Y. Cheng, J. Zhu, D. Xue, G. Abt, H. Ye, and Y. Peng, "A novel adaptive spectrum noise cancellation approach for enhancing heartbeat rate monitoring in a wearable device," *IEEE Access*, vol. 6, pp. 8364–8375, 2018.
- [8] T. Schäck, M. Muma, and A. M. Zoubir, "Computationally efficient heart rate estimation during physical exercise using photoplethysmographic signals," in *2017 25th European Signal Processing Conference (EUSIPCO)*. IEEE, 2017, pp. 2478–2481.
- [9] S. M. Salehizadeh, D. Dao, J. Bolkhovskiy, C. Cho, Y. Mendelson, and K. H. Chon, "A novel time-varying spectral filtering algorithm for reconstruction of motion artifact corrupted heart rate signals during intense physical activities using a wearable photoplethysmogram sensor," *Sensors*, vol. 16, no. 1, p. 10, 2015.
- [10] N. Huang and N. Selvaraj, "Robust ppg-based ambulatory heart rate tracking algorithm," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2020, pp. 5929–5934.
- [11] M. Zhou and N. Selvaraj, "Heart rate monitoring using sparse spectral curve tracing," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2020, pp. 5347–5352.
- [12] A. Reiss, I. Indlekofer, P. Schmidt, and K. Van Laerhoven, "Deep ppg: Large-scale heart rate estimation with convolutional neural networks," *Sensors*, vol. 19, no. 14, p. 3079, 2019.
- [13] P. Kasnesis, L. Toumanidis, A. Burrello, C. Chatzigeorgiou, and C. Z. Patrikakis, "Multi-head cross-attentional ppg and motion signal fusion for heart rate estimation," *arXiv preprint arXiv:2210.11415*, 2022.
- [14] A. Burrello, D. J. Pagliari, M. Risso, S. Benatti, E. Macii, L. Benini, and M. Poncino, "Q-ppg: Energy-efficient ppg-based heart rate monitoring on wearable devices," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 15, no. 6, pp. 1196–1209, 2021.
- [15] V. Bieri, P. Strelti, B. U. Demirel, and C. Holz, "Beliefppg: Uncertainty-aware heart rate estimation from ppg signals via belief propagation," *arXiv preprint arXiv:2306.07730*, 2023.
- [16] D. Ray, T. Collins, and P. V. Ponnappalli, "Deeppulse: An uncertainty-aware deep neural network for heart rate estimations from wrist-worn photoplethysmography," in *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2022, pp. 1651–1654.
- [17] A. Burrello, D. J. Pagliari, M. Bianco, E. Macii, L. Benini, M. Poncino, and S. Benatti, "Improving ppg-based heart-rate monitoring with synthetically generated data," in *2022 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2022, pp. 153–157.
- [18] L. Von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, J. Frommer, A. Pick, R. Ramamurthy *et al.*, "Informed machine learning—a taxonomy and survey of integrating prior knowledge into learning systems," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 1, pp. 614–633, 2021.
- [19] G. Masinelli, F. Dell'Agnola, A. A. Valdés, and D. Aienza, "Spare: A spectral peak recovery algorithm for ppg signals pulsewave reconstruction in multimodal wearable devices," *Sensors*, vol. 21, no. 8, p. 2725, 2021.
- [20] S. B. Song, J. W. Nam, and J. H. Kim, "Nas-ppg: Ppg-based heart rate estimation using neural architecture search," *IEEE Sensors Journal*, vol. 21, no. 13, pp. 14941–14949, 2021.
- [21] M. Risso, A. Burrello, D. J. Pagliari, S. Benatti, E. Macii, L. Benini, and M. Pontino, "Robust and energy-efficient ppg-based heart-rate monitoring," in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2021, pp. 1–5.
- [22] L. Scimeca, S. J. Oh, S. Chun, M. Poli, and S. Yun, "Which shortcut cues will dnns choose? a study from the parameter-space perspective," *arXiv preprint arXiv:2110.03095*, 2021.
- [23] J. Y. A. Foo and S. J. Wilson, "A computational system to optimise noise rejection in photoplethysmography signals during motion or poor perfusion states," *Medical and Biological Engineering and Computing*, vol. 44, pp. 140–145, 2006.
- [24] Y. Ye, Y. Cheng, W. He, M. Hou, and Z. Zhang, "Combining nonlinear adaptive filtering and signal decomposition for motion artifact removal in wearable photoplethysmography," *IEEE Sensors Journal*, vol. 16, no. 19, pp. 7133–7141, 2016.
- [25] S. Kim, S. Im, and T. Park, "Characterization of quadratic nonlinearity between motion artifact and acceleration data and its application to heartbeat rate estimation," *Sensors*, vol. 17, no. 8, p. 1872, 2017.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" *Advances in neural information processing systems*, vol. 30, 2017.
- [29] J. Cho, Y. Sung, K. Shin, D. Jung, Y. Kim, and N. Kim, "A preliminary study on photoplethysmogram (ppg) signal analysis for reduction of motion artifact in frequency domain," in *2012 IEEE-EMBS Conference on Biomedical Engineering and Sciences*. IEEE, 2012, pp. 28–33.
- [30] P. R. Rijnbeek, G. Van Herpen, M. L. Bots, S. Man, N. Verweij, A. Hofman, H. Hillege, M. E. Numans, C. A. Swenne, J. C. Witteman *et al.*, "Normal values of the electrocardiogram for ages 16–90 years," *Journal of electrocardiology*, vol. 47, no. 6, pp. 914–921, 2014.
- [31] L. Chhabra, N. Goel, L. Prajapat, D. H. Spodick, and S. Goyal, "Mouse heart rate in a human: diagnostic mystery of an extreme tachyarrhythmia," *Indian pacing and electrophysiology journal*, vol. 12, no. 1, pp. 32–35, 2012.