




# Nightbeat: Heart Rate Estimation From a Wrist-Worn Accelerometer During Sleep

Max Moebus , Graduate Student Member, IEEE, Lars Hauptmann, Nicolas Kopp, Berken Utku Demirel , Björn Braun, and Christian Holz , Member, IEEE

**Abstract**—Today’s fitness bands and smartwatches typically track heart rates (HR) using optical sensors. Large behavioral studies such as the U.K. Biobank use activity trackers without such optical sensors and thus lack HR data, which could reveal valuable health trends for the wider population. In this paper, we present the first dataset of wrist-worn accelerometer recordings and electrocardiogram references in uncontrolled at-home settings to investigate the recent promise of IMU-only HR estimation via ballistocardiograms. Our recordings are from 42 patients during the night, totaling 310 hours. We also introduce a frequency-based method to extract HR via curve tracing from IMU recordings while rejecting motion artifacts. Using our dataset, we analyze existing baselines and show that our method achieves a mean absolute error of 0.88 bpm—76% better than previous approaches and the first to surpass established medical standards for heart rate monitors. Our results validate the potential of IMU-only HR estimation as a key indicator of cardiac activity in existing longitudinal studies to discover novel health insights.

**Index Terms**—Accelerometers, heart rate, sleep.

## I. INTRODUCTION

WRIST-WORN accelerometers are commonly used in longitudinal studies to analyze patients’ activity patterns, such as exercise routines, physical activity metrics, or sleep and wake times. For example, the U.K. Biobank combines wrist-based accelerometer recordings from 100,000 patients during a week-long sub-study [1]. This study alone has revealed various activity and sleep patterns as risk factors for Parkinson’s [2], depression [3], cardiovascular disease [4], different types of cancer [5], and overall mortality [6]. Besides the accelerometer, no other sensing modality was included. Because the U.K. Biobank and comparable studies [7], [8] pair

week-long accelerometer recordings with decades of electronic health records, using the accelerometer recordings to extract physiological features that are otherwise not included in the study could advance early disease detection and risk factor identification on a population level. Heart rate dynamics have been shown to link to sleep quality and the progression of rare diseases alike, hence, showing great value for identifying novel risk factors on a population level [9], [10], [11].

Recent studies have shown promise for analyzing cardiac activity from the data recorded by body-worn accelerometers, specifically heart rate (HR [12], [13]). These approaches leverage the ballistocardiogram (BCG [14], [15]), which captures the subtle mechanical vibrations caused by heartbeats.

Extracting robust HR estimates from BCG recordings in real-world settings is a substantial challenge, particularly when sampled from the wrist. Even simple daytime activities such as walking or eating exhibit larger motion magnitudes than those in the BCG [12], [14], rendering such estimation near impossible in practical settings. Moments of sleep mark a notable exception—they lend themselves to sampling BCG signals, as voluntary movements are minimal [12], [13], [14]. Previous studies have extracted HR values from individual patients’ recordings in controlled sleep laboratories [12], [13], [14], where sensor placement was verified by staff. Here, even individual heartbeats can be detected from wrist-worn accelerometers during sleep [12], as long as motion-afflicted segments are reliably removed (e.g., up to 80% of recordings [13]). Despite this promise, no existing methods reach practical accuracies in their estimates for uncontrolled real-world environments—and no datasets with continuous ground-truth references exist to facilitate and evaluate developments of the former.

In this paper, we introduce Nightbeat-DB, a novel dataset for BCG-based HR estimation tasks during sleep. Our dataset replicates the settings of large longitudinal studies such as the U.K. Biobank [1] to allow for the evaluation of techniques in similarly uncontrolled environments. 42 patients received a wrist-worn activity tracker (Axivity AX3, which embeds a 3-axis accelerometer) for wear on their dominant wrist for one night at home while following their regular sleep routine at home. Patients also received an ECG chest belt (movisens EcgMove 4) to record reference signals for ground-truth values of HR and inter-beat intervals (IBI). Patients also indicated if they shared their beds or slept alone, which is important for BCG analyses, as vibrations couple through bed mattresses and can lead to interfering signals [16].

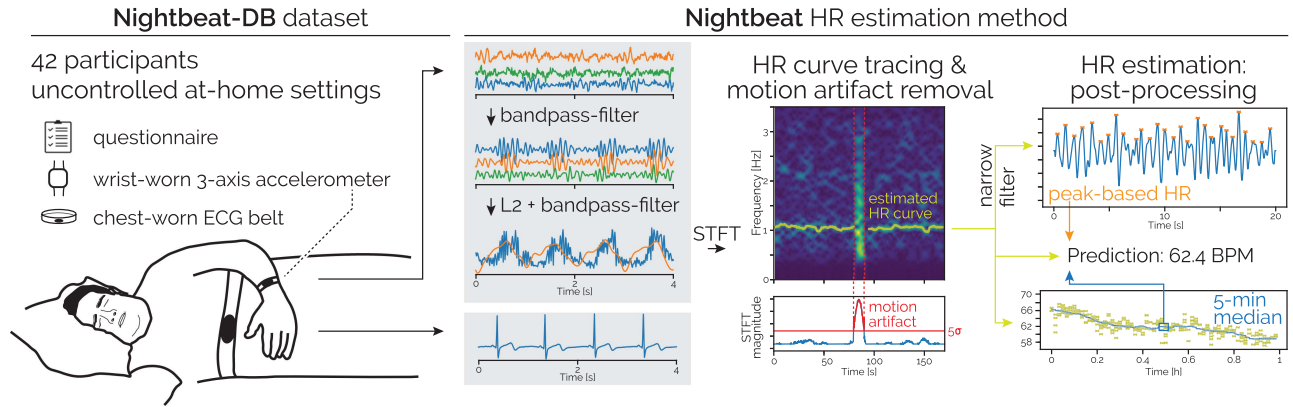
Received 24 November 2024; accepted 29 November 2024. Date of publication 9 December 2024; date of current version 6 February 2025. (Corresponding author: Christian Holz.)

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by Ethics Commission of ETH Zurich (under Application No. 23 ETHICS-002, and performed in line with the Declaration of Helsinki).

The authors are with the Department of Computer Science, ETH Zurich, 8092 Zürich, Switzerland (e-mail: max.moebus@inf.ethz.ch; lars.hauptmann@inf.ethz.ch; nicolas.kopp@inf.ethz.ch; berken.demirel@inf.ethz.ch; bjoern.braun@inf.ethz.ch; christian.holz@inf.ethz.ch).

Our dataset, Nightbeat-DB, and our source code are available on GitHub: <https://github.com/eth-siplab/Nightbeat>.

Digital Object Identifier 10.1109/JBHI.2024.3514293



**Fig. 1.** Our dataset comprises continuous motion signals from a wrist-worn 3-axis accelerometer (Axivity AX3, 100 Hz) and corresponding ECG signals (movisens EcgMove 4) from 42 patients and nights. From the ballistocardiogram (BCG) captured by the vibrations reaching the wrist-based sensor, our signal processing method estimates the patient’s heart rate using a combination of filters and heuristics for pre-processing, a short-term frequency analysis to identify HR curve, and a filter stage to precisely select the peak of the BCG wave for inter-beat interval detection. We combine the HR curve with the detected inter-beat intervals and a 5-minute median smooth to make a prediction for every 20-second window.

We complement the dataset with Nightbeat, an estimation method for HR from 3-axis wrist-based accelerometer recordings by removing motion artifacts, tracing HR curves in the frequency domain, detecting individual heartbeats after a filter and frequency detection stage, followed by simple post-processing. We validate our method on Nightbeat-DB and show its robustness to intermittent motion during the night. Compared to 3 baseline methods, our method achieves the lowest error ( $MAE = 0.88$  bpm), lowering the error of previous approaches ( $MAE = 3.68$ ) by 76%, while removing the same amount of data (22%) [12]. For female participants, we find that sharing a bed leads to an average increase in MAE of 32%. For male participants, the difference is only 2%.

In summary, this paper makes the following contributions:

- 1) a novel dataset of continuous wrist-based accelerometer signals and paired ECG reference signals during the night for HR estimation tasks, designed to replicate the uncontrolled at-home setting of existing large-scale datasets that do not include continuous HR recordings “in-the-wild” [1], [7], [8],
- 2) a signal processing method for HR estimation from accelerometer signals that rejects motion artifacts, detects the HR via curve tracing in the frequency domain, and estimates heartbeats and inter-beat intervals, and
- 3) a comparison of existing approaches for HR estimation from wrist-worn accelerometers using BCG sensations on our dataset in the context of uncontrolled at-home settings where some patients share their beds. Our method achieves the lowest error compared to 3 (recent) baseline methods ( $MAE = 0.88$ , 76% reduction compared to recent baselines) and an average correlation of 0.81 across participants.

Taken together, we demonstrate the suitability of our method for application to existing and future datasets collected in free-living and at-home conditions with activity monitors that do not include alternative sensors for HR estimation (e.g., U.K. Biobank). Our method achieves accuracy that is suitable for

medical analysis [12], making it valuable for future activity studies and monitoring cardiac activity in natural settings.

## II. RELATED WORK

### A. HR Estimation From Wearable Sensors

Most modern wearables use optical sensors for HR estimation. While highly successful in periods with little motion, HR estimation based on Photoplethysmography (PPG) becomes difficult in the presence of motion artifacts [14]. While a recent method fuses PPG and accelerometer signals in a probabilistic deep learning framework to deal with motion artifacts [17], most works still employ simple rules on aggregates of the underlying signal or the short-time Fourier transform (STFT) to identify segments corrupted by motion artifacts [18]. After reliably detecting corrupted segments, HR estimation on the remaining uncorrupted signals proves a much simpler task where well-designed signal processing methods often outperform (convolutional) neural networks [18].

### B. Signal Quality Estimation

Estimating the quality of a signal segment (e.g. detecting motion artifacts) has been crucial in health monitoring across various sensing modalities [19], [20], [21]. Most of these methods extract several features, so-called signal quality indices (SQI), from the signals and employ rule-based decision algorithms to detect segments with artifacts [20]. Alternatively, agreement in prediction between different sensors or template matching have been used to identify segments of high signal quality [20], [21]. Although some learning-based algorithms have been proposed [19], they require ground truth labels, which must be annotated manually by an expert. In this work, we propose a method for motion artifact detection based on the energy in the STFT that does not require ground truth annotations for training. We further use disagreement between different estimation techniques to identify potentially corrupted segments.

### C. Datasets for Accelerometer-Based HR Estimation During Sleep

Accelerometers have a decades-long history in sleep science [22], yet only few datasets are suitable for HR estimation from wrist-worn accelerometers during sleep. Besides sampling at a sufficient rate, accelerometers must also be sensitive enough to detect the subtle BCG waveforms. Currently available datasets that allow to detect BCG waveforms, however, have been recorded in sleep laboratories, where signal quality is higher than in uncontrolled at-home settings.

Multiple datasets are suitable for our task of HR estimation from wrist-worn accelerometers [12], [14], [23]. Walch et al. recorded a dataset with 32 participants for sleep staging based on signals supplied by the Apple Watch, i.e. a 3-axis accelerometer sampling at 50 Hz and HR values supplied by the Apple Watch's PPG sensor every 5–15 seconds [23], [24]. For the task of sleep staging, the dataset was recorded in a sleep laboratory. Also inside a sleep laboratory, Zschocke et al. recorded a dataset using the SOMNO-Watch comprising of more than 200 participants to reconstruct respiration and HR signals [13], [14]. Similarly, Weaver et al. collected data from 82 children inside a sleep laboratory and compared the performance of an Apple Watch and ActiGraph GT9X for accelerometer-based HR estimation [12].

Current datasets have been recorded in sleep laboratories [12], [14], [23], which limits their use for the purpose of HR estimation from wrist-worn accelerometers in uncontrolled at-home settings. Sleep laboratories are not representative of at-home settings, mainly since sleep laboratories are limited to participants sleeping alone in their beds and thus omit a potentially crucial source of noise. Motions emitted during the cardiac cycle are measurable through a mattress [16] leading to potential interference if two people share a bed. Furthermore, in large studies such as the U.K. Biobank, participants receive the accelerometer via mail [1] and are instructed to put it on themselves. In sleep laboratories, trained personnel verify that sensors are worn correctly and ensure signal quality.

### D. Accelerometers for HR Estimation During Sleep

Accelerometers have been used in diverse settings for HR monitoring, even though not currently when worn on the wrist in completely uncontrolled settings. Seismocardiography, when an accelerometer is placed on the chest, is likely the most common setting [25]. Without motion artifacts, it is then possible to observe the characteristic BCG waveform.

The characteristic BCG components are also visible if an accelerometer is integrated into a chair [26] or placed under a scale [27] that the participant is standing or sitting on. During sleep, accelerometers have been successful at HR as well as breathing rate estimation if placed on or under a mattress [16], [28]. For wrist-worn accelerometers, multiple approaches have been proposed for extracting HR and heartbeats during sleep [12], [14], [15]. For example, pulse wave reconstruction (PWR) [13], [14] applies a bandpass filter with cut-off frequencies of 5–14 Hz followed by a Hilbert transform to detect peaks on the axis with the highest autocorrelation in the plausible range of HR values. The authors also exclude segments where the total magnitude is above 5 mg and segment

the accelerometer signal based on ECG peaks to demonstrate correlations of inter-beat intervals of up to 0.91. On average they ignore everything apart from 1.7 h per night per participant. The requirement for ground truth ECG peaks makes the approach impractical “in-the-wild”. Similarly, BioInsights [15] requires completely movement-free periods for HR extraction from wrist-worn accelerometers without presenting a technique to remove motion artifacts. Even during low-movement periods such as asleep, this is impractical. As a result, both PWA and BioInsights highlight the need for more robust methods that can accurately extract HR data despite the presence of motion artifacts. A recent study by Weaver et al. investigated extracting HR from accelerometers in more realistic sleep settings for children [12]. The authors proposed calculating the derivative of the total magnitude of the wrist-worn accelerometer, followed by root mean square averaging, and then applying the Hilbert transform while using a short-time Fourier transform for HR detection. Using the accelerometer signal of an Apple Watch (sampled at 50 Hz), they achieve a mean absolute error of 6.4 bpm, which increases to 16.8 bpm when using an ActiGraph GT9X (sampled at 100 Hz). The authors use the rolling standard deviation of the STFT peak-width (140 seconds) to estimate signal quality.

In this work, we present Nightbeat-DB, a comprehensive dataset recorded in at-home settings, for a more robust evaluation of HR extraction methods. Further, we present Nightbeat, a HR estimation method capable of detecting motion artifacts and accurately tracking the HR across the entire night.

## III. NIGHTBEAT-DB DATASET

To study the performance of HR estimation tasks based on wrist-worn accelerometers during sleep outside sleep laboratories, we recorded a dataset with 42 patients “in the wild.” In contrast to existing datasets, the characteristics of our dataset mimic the setting of existing large-scale health studies that include wrist-worn accelerometers, yet no HR signals.

### A. Study Setup and Participants

All participants received a 3-axis accelerometer for wear on their dominant wrist (Axivity AX3) and an ECG chest belt (movisens EcgMove 4). Participants received the same simple and short instructions as those in the U.K. Biobank motion substudy [1] to ensure a little-controlled experimental setting. Participants were instructed to wear the accelerometer such that the z-axis was orthogonal to the skin surface and that the wristband did not “jiggle” yet was still comfortable for a single night: from shortly before going to bed until shortly after waking up. Similarly, participants were instructed to tightly wear the ECG belt, but not to tighten it beyond comfort. Further, participants were asked to fill in a questionnaire after waking up about demographic information, the time they went to bed and woke up, and if they slept alone. In the case they did not sleep alone in their bed, participants indicated whether they slept on separate mattresses.

All 42 participants provided written informed consent and the study protocol was approved by the ethics committee of ETH Zurich: 23 ETHICS-002. 42 participants were recruited

via convenience sampling in and around Zurich (18 female, 24 male, ages 20–36,  $M = 26.48$  years). None reported to suffer from a sleep disorder. 15 participants reported that they shared their bed with someone else when participating in the study. None of these 15 participants slept on a separate mattress.

### B. Recorded Signals and Sensor Synchronization

The Axivity AX3 recorded motion signals at the dominant wrist at 100 Hz ( $\pm 8$  g). The movisens EcgMove 4 recorded a 1-channel ECG signal at 1024 Hz. The movisens EcgMove 4 further recorded motion from the chest at 64 Hz ( $\pm 16$  g), which we used to synchronize the two sensors. Before participants were handed the two devices, we initiated the recording and included a synchronization pattern by laying the two devices on top of one another and hitting them against a book several times. After we received back the two devices, we included another pattern for synchronization and stopped the recording. To synchronize the signals from the two devices, we upsampled the motion signal from the movisens EcgMove 4, manually spotted the approximate time stamps of the synchronization patterns, and then used the cross-correlation between the two motion signals to accurately align them (motivated by the procedure described in [29]).

### C. Ground-Truth HR Measurements From ECG

We extracted reference HR values as ground truth from ECG recordings in two ways. First, we used the vendor’s built-in algorithm to obtain continuous HR measurements (movisens EcgMove 4). Second, we additionally processed the ECG recordings with a series of methods implemented as part of the NeuroKit2 Python library [30], including popular algorithms such as Pan-Tompkins [31] but also a more modern wavelet-based approach [32]. We compared the output from all algorithms with that of the vendor’s implementation and manually verified each signal portion for disagreements in HR estimates. Since no inspected signal recordings were of low quality, we selected HR estimates from the vendor. For a detailed inspection, please visit our GitHub repository.

### D. Data Preparation & Format

We made use of the approximate sleep and wake times provided by participants to set the start and end points for our processing pipeline. Subsequently, we used the provided signal quality metric by the movisens EcgMove 4 (provided at 1/60 Hz) to identify the first 60-second window after the sleep time provided by participants when the participant wore the movisens EcgMove 4. From there until the last 60-second window when participants wore the movisens EcgMove 4 before their approximate wake time, we generally processed the synchronized raw data in 3-minute windows with a 2-minute overlap. 3-minute processing windows proved long enough such that our method of motion artifact detection and HR curve tracing worked well (Sections IV-B and IV-C). For the central 60-second window, we primarily relied on the data quality metric provided by the movisens EcgMove 4 to detect potential non-wear, ensure the

participants wore the ECG belt correctly such that a meaningful ECG signal was measurable, and that there were no motion artifacts that affected at least 5 consecutive heartbeats. As part of our proposed approach, we removed further motion artifacts as outlined in Section IV-B.

## IV. NIGHTBEAT HR ESTIMATION METHOD

Our proposed method estimates HR from wrist-worn accelerometers by detecting the BCG waveforms in the accelerometer signal. To make the BCG waveforms visible, we calculate the L2-norm of the three bandpass-filtered accelerometer axes and refer to this transformed signal as the Nightbeat signal (Section IV-A). After computing the Nightbeat signal, our algorithm consists of four main parts:

- 1) Motion artifact removal (Section IV-B)
- 2) HR curve tracing based on the STFT (Section IV-C),
- 3) Heartbeat detection in the time domain on a narrowly filtered Nightbeat signal (Section IV-D), and
- 4) Robust post-processing (Section IV-E).

After transformation, the Nightbeat signal is split into 3-minute windows with a 2-minute overlap to identify motion artifacts and to detect the overall trajectory of participants’ HR. Our method for HR estimation is inspired by curve tracing for PPG [18] and tracks peaks in the STFT output throughout each 3-minute window. For 20-second windows with a 10-second overlap that are free of motion artifacts, we take the mean value of the estimated HR trajectory (or HR curve) as our frequency-based HR prediction ( $HR_C$ ). We then narrowly filter the Nightbeat signal of each 20-second window around the estimated HR frequency  $HR_C$  to detect heartbeats via peak detection in the time domain. From the detected inter-beat intervals, we then derive a peak-based HR prediction ( $HR_P$ ). Finally, for all windows where  $|HR_C - HR_P| \leq 10$ , we compute the rolling median of the frequency-based HR predictions across 5-minute windows ( $HR_M$ ). If for any window,  $|HR_C - HR_P| > 10$  or  $|HR_C - HR_E| > 10$ , we discard the window. Else, we predict  $HR_C$ .

### A. Nightbeat Signal: Filtering & L2-Norm

To arrive at the Nightbeat signal, we first filter each of the three axes using a 4<sup>th</sup>-order Butterworth bandpass filter to isolate frequencies in the range of 5–14 Hz and compute the L2-norm of the 3 filtered axes. We then band-pass filter it using a 4<sup>th</sup>-order Butterworth filter to the range of 0.5 to 3.5 Hz. The procedure is visualized in Fig. 1.

### B. Artifact Removal

Abrupt or excessive motion obscures the BCG waveforms in the Nightbeat signal. While such movements occur less frequently during sleep than while awake, the Nightbeat signal is not free of such artifacts. We detect motion artifacts directly from the time-frequency representation (STFT) of the 3-minute Nightbeat signal. To do so, we compute the STFT with a window length of 1024 samples (Kaiser window function with  $\beta = 2$ ) padding each window by factor 10 and shifting it by 10 samples.

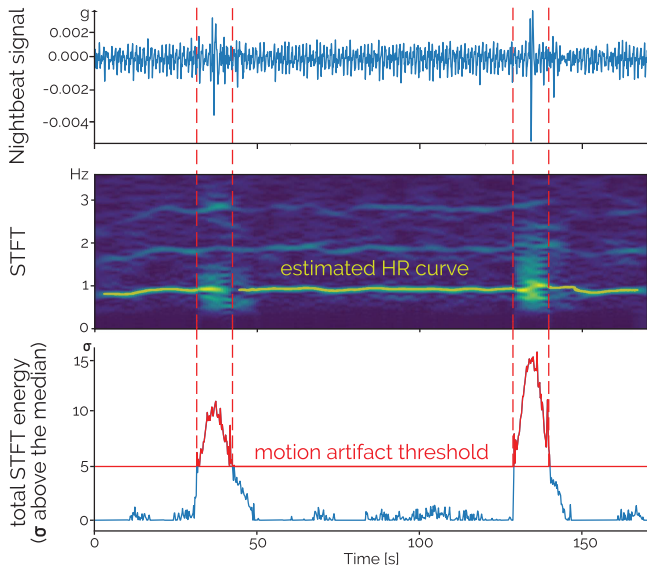


Fig. 2. We detect motion artifacts during 3-minute windows based on the energy of the STFT of the Nightbeat signal over time. If the energy lies 5 standard deviations ( $\sigma$ ; estimated robustly via the IQR) above the median, we flag the segment as a motion artifact.

The output of the STFT had a resolution of 0.2 seconds and 0.01 Hz, in the time and frequency dimensions, respectively.

From the STFT of the 3-minute window, we sum the magnitude across all frequency bins for each timestep and essentially derive a distribution of total spectral energy over time. As shown in Fig. 2, we identify segments as motion artifacts where the total spectral energy lies 5 standard deviations ( $\sigma$ ) above the median of the 3-minute window. Window lengths of 3 minutes ensured that motion artifacts could be reliably identified. We estimate  $\sigma$  robustly using the interquartile range (IQR; the IQR of  $N(\mu, \sigma)$  is  $1.349\sigma$ ).

### C. Frequency Peak Identification: Curve Tracing

Based on the same STFT output used for artifact removal in Section IV-B (i.e., STFT of the 3-minute Nightbeat signal), we estimate HR values via curve tracing during periods free of motion artifacts [18]. While our method is inspired by [18], we do not use gradients to identify curves but simply track possible peaks in the frequency domain across consecutive timesteps—allowing for small gaps and jumps.

At each timestep (every 0.2 seconds) of the STFT output, we identify all peaks in the frequency domain (in the range 0.5–3.5 Hz) whose height and prominence exceed half of the highest magnitude in the range of 0.5–3.5 Hz. As visible in Fig. 2, BCG harmonics are often visible at multiples of the HR frequency. To avoid tracking the BCG harmonics, we ignore any peaks that occur at a multiple of the frequency of a lower-frequency peak (up to a deviation of 2%). After identifying all valid peaks at every timestep, we connect neighboring peaks to what we refer to as ‘HR curves’. Starting at the first timestep, all identified peaks in the frequency domain start a new HR curve. Moving to the next timestep, all identified peaks are attributed to an existing HR curve if they lie within 0.05 Hz of a peak at the previous




Fig. 3. Peak detection (purple crosses) from the Nightbeat signal (blue) for an exemplary 15-second signal. The smoothed and filtered signal is displayed in orange, the height threshold for peak detection in green.

timestep or start a new HR curve. If for any identified peak there was no peak at the previous timestep within 0.05 Hz we looked up to 5 timesteps back—effectively allowing for 1 s gaps in HR curves. Any peak that cannot be connected to an already existing HR curve starts a new HR curve. After attributing all peaks to a curve, we then remove all curves that last less than 10 seconds. We are then left with potentially multiple curves at the same timestep. We proceed by calculating the median frequency of all points attributed to any curve that is longer than 10 seconds and their IQR. We then ignore any curve whose median is further than  $5\sigma$  away from the overall median—where the  $\sigma$  is again robustly estimated via the IQR. In the case that there are two curves remaining for the same timestep, we select the curve closest to the overall median. We then arrive at curves such as the ones displayed in Figs. 1 and 2, where there is at most one HR curve per timestep but the curves might have gaps. We interrupt any HR curves during motion artifacts and start anew after the identified motion artifact.

For each 20-second window, we then arrive at a frequency-based HR prediction based on the mean of the detected HR curve during this window:  $HR_C$  (in Hz). We skip any window where no HR curve has been detected.

### D. Heart Beat Detection

Based on  $HR_C$  (in Hz) for a 20-second window, we perform peak detection in the time domain on a signal narrowly filtered around  $HR_C$ . First, we apply the rolling median across  $0.3 s/HR_C$ , followed by taking the rolling mean across the same window length to lowpass filter and smooth the signal. We then narrowly bandpass filter this signal to the range of  $0.9HR_C - 1.1HR_C$  using a 20<sup>th</sup>-order FIR filter. We then detect all peaks with a height exceeding a rolling mean across  $0.5 s/HR_C$  (Fig. 3). We then average the inter-beat intervals to arrive at our second HR prediction:  $HR_P$ . If  $|HR_C - HR_P| > 10$  bpm, we do not make a prediction. In our GitHub repository, we further compare alternative approaches.

### E. Post-Processing: Rolling Median

After ignoring 20-second windows, where  $|HR_C - HR_P| > 10$  bpm, we compute  $HR_M$ —the rolling 5-minute median of  $HR_C$ . If  $|HR_C - HR_M| \leq 10$  bpm, we predict  $HR_C$ , else no prediction is made. The robustness of the rolling median reliably detects outlier predictions yet is still able to follow heart rate dynamics during sleep.

## V. VALIDATION

We evaluate Nightbeat on Nightbeat-DB and the AW dataset (Section V-A) against 3 baselines (Section V-B). Together with Nightbeat-DB, and our Nightbeat method, all relevant code to replicate our analysis is available via GitHub.

### A. AW Dataset

In addition to Nightbeat-DB, we used a dataset recorded with an Apple Watch (Series 2 & 3) that comprises 31 participants (21 female) where motion data from a 3-axis accelerometer was recorded at 50 Hz from the wrist [23], [33]. HR values are available approximately every 5 seconds based on the Apple Watch’s PPG sensor. Initially, the dataset was recorded to assess sleep staging algorithms. Participants were aged 19–55 years (29.4 years on average) and the available data was recorded during the participants’ stay in a sleep laboratory. Only HR estimates supplied by the Apple Watch are available. We will refer to this dataset simply as the ‘AW’ dataset. It is publicly available via Physionet [24].

We processed the AW dataset similarly to our Nightbeat-DB dataset (Section III-D) based on 3-minute windows with a 2-minute overlap. However, we used different data requirements adapted to the AW dataset. The Apple Watch returned slightly irregularly sampled acceleration signals. Due to the specifications of our algorithm as outlined in Section IV, we required that no gap between consecutive samples exceeded 0.05 seconds and that no gap between consecutive HR values exceeded 10 seconds. If these requirements were met, we upsampled the output of the accelerometer signal to 100 Hz via linear interpolation to make the sampling rate match our Nightbeat-DB dataset. In total, the data amounted to 113 hours.

### B. Implementations of Alternative Approaches

In addition to Nightbeat outlined in Sections IV-A–IV-E, we compare three approaches from related works described in Section II-D: BioInsights by Hernandez et al. [15], Pulse Wave Reconstruction (PWR) by Zschocke et al. [13], [14] and Jerks are Useful (Jerks) by Weaver et al. [12]. To the best of our knowledge, these three approaches summarize the current works for HR estimation from wrist-worn accelerometers while asleep. If applicable, we implemented the signal quality metrics proposed by the respective approach and removed the same percentage of 20-second windows according to this metric as removed by Nightbeat (approx. 22%) to allow for fair comparison. However, Zschocke et al. also remove beats during their evaluation if they are not detected within a specified window following an ECG peak [13]. Since we want the methods to be completely independent of any signal not derived from the accelerometer, we did not implement this last quality check. We reimplemented all approaches in Python and they can be found in our Github repository.

Besides the three approaches described above, we implemented a personalized Mean-Baseline (Base: Mean-Pred) that simply predicts the average HR label for each participant.

TABLE I

COMPARISON OF DIFFERENT ALGORITHMS FOR HR PREDICTION FROM WRIST-WORN ACCELEROMETERS WHILE ASLEEP

approach	dataset	MAE	RMSE	Cor
BioInsights [15]	Nightbeat-DB	17.47	21.38	0.00
	AW	21.12	24.53	-0.03
Jerks [12]	Nightbeat-DB	13.23	20.63	0.04
	AW	3.68	7.11	0.50
PWR [13]	Nightbeat-DB	8.83	10.91	0.22
	AW	6.05	7.77	0.37
Base: Mean-Pred.	Nightbeat-DB	3.16	4.07	0
	AW	2.98	4.09	0
<b>Nightbeat (ours)</b>	Nightbeat-DB	<b>0.88</b>	<b>2.24</b>	<b>0.81</b>
	AW	<b>1.68</b>	<b>3.38</b>	<b>0.64</b>




Fig. 4. Comparison of Prediction and ground-truth HR over time for P11 of the Nightbeat-DB dataset. Using our Nightbeat HR estimation method, we achieve an MAE of 0.41 and a correlation of 0.98.




Fig. 5. Predictions against ground-truth HR for all participants of the Nightbeat-DB dataset—colored by participants. Across all participants, our Nightbeat HR estimation method achieves a correlation of 0.94.

## VI. RESULTS

Table I shows the results of Nightbeat against three baseline approaches. We use the mean absolute error (MAE), root mean squared error (RMSE), and correlation (Cor) as metrics to compare the different approaches across the two datasets.

Nightbeat outperforms all Baselines and achieves an MAE of as low as 0.88 bpm averaged across all participants on Nightbeat-DB. Averaged across individual participants, we achieve a correlation of 0.81. For a single participant, we achieve a correlation as high as 0.98 on Nightbeat-DB at an MAE of 0.41 bpm (Fig. 4). Across all participants, the correlation exceeds 0.95 (Fig. 5).

**TABLE II**  
PERFORMANCE OF NIGHTBEAT DEPENDING ON SEX AND WHETHER PARTICIPANTS SLEPT ALONE (NIGHTBEAT-DB)

approach	sex	alone	MAE	RMSE	Cor
Nightbeat (ours)	F	Y	0.78	1.89	0.87
	F	N	1.03	2.30	0.80
	M	Y	0.86	2.49	0.77
	M	N	0.88	2.05	0.84
	All	Y	0.83	2.23	0.82
	All	N	0.94	2.15	0.82
Base: Mean-Pred.	F	Y	3.01	3.95	0
	F	N	2.71	3.47	0
	M	Y	3.19	4.28	0
	M	N	3.01	3.86	0
	All	Y	3.11	4.14	0
	All	N	2.88	3.70	0

Out of all possible 20-second windows, we remove 22% due to estimated motion artifacts, gaps in HR curves, or due to the  $HR_C$  prediction lying further than 10 bpm away from either of the  $HR_P$  or  $HR_M$  prediction. On the AW dataset, we achieve an MAE of 1.68 bpm at an average correlation of 0.64. Nightbeat is the only method that achieves lower errors than when predicting an individual’s average HR during the entire night.

PWR and Jerks achieve MAEs of 8.83 and 13.23, respectively on the Nightbeat-DB. On the AW dataset, PWR and Jerks achieve lower MAEs of 6.05 and 3.68 bpm, respectively, at average correlations of 0.37 and 0.50. The BioInsights approach achieves high MAEs of around 20 bpm on both datasets at a correlation of approximately 0. None of the three approaches achieves MAEs below 3.36 or 2.98 bpm, which is the performance of predicting each participant’s average HR on Nightbeat-DB and the AW dataset (Base: Mean-Pred.).

In Table II, we notice a 32% increase in MAE for female participants when sharing a bed (based on Nightbeat). For male participants, the MAE does not increase when sharing a bed. In fact, it decreases by 2%. On Nightbeat-DB, MAE was not significantly correlated with age ( $r = 0.18, p = 0.27$ ), weight ( $r = 0.03, p = 0.84$ ), size ( $r = -0.05, p = 0.77$ ), and average accelerometer magnitude ( $r = -0.20, p = 0.24$ ).

## VII. DISCUSSION

With Nightbeat and Nightbeat-DB, we present a novel dataset that closely replicates the setting of large longitudinal studies accompanied by a novel approach that is suitable for HR prediction in such uncontrolled environments. Various large longitudinal studies comprise week-long accelerometer recordings and decades of health records that are constantly being updated, yet lack continuous HR signals “in-the-wild”. Using continuous HR signals (while asleep), we might advance our understanding of diseases such as cardiovascular disease, mental disorders, or even neurological diseases and their respective disease progression and risk factors. Compared to (recent) baselines, Nightbeat is the only approach that exceeds mean prediction in terms of MAE during the night and meets well-established medical standards for HR monitors set by the FDA with a reduction in

errors by at least 76% across datasets. Regulatory benchmarks for cardiac monitors, such as those used by the FDA, allow an error margin of  $\pm 10\%$  or  $\pm 5$  bpm (whichever is greater) and our method is the only one to fall well within these thresholds—even in terms of RMSE [34], [35]. We thus take a large step towards HR estimation while asleep in uncontrolled studies such as the U.K. Biobank accelerometer study that currently lack HR signals for medical analysis.

Nightbeat performs almost twice as well on Nightbeat-DB, despite Nightbeat-DB a less controlled environment. To derive ground-truth HR values for the AW dataset, we essentially linearly interpolate the labels provided every 5–15 seconds by the Apple Watch. This likely leads to inaccuracies in the labels and a glass ceiling for performance.

In contrast to Nightbeat, PWR and Jerks perform better on the AW dataset. This might be indicative of the more challenging environment of at-home settings compared to sleep laboratories. Still, the increase in MAE from 3.68 to 13.23 for Jerks is surprising. In their original manuscript, however, Weaver et al. observed a similar pattern and reported an MAE of 6.4 when using the Apple Watch, but an MAE of 16.8 when using the GT9X. This might indicate that their jerk-based approach is sensitive to the specific accelerometer. As arguably the simplest baseline and without motion artifact removal, it is unsurprising that Bioinsights does not cope well in realistic scenarios. While significantly more sophisticated than BioInsights, PWR seems optimized for segments with very high signal quality, which might explain its relatively low performance if only removing 22% of recordings.

Nightbeat combines a relatively simple, yet robust signal aggregation, with a sophisticated approach for HR estimation that makes use of the relatively slow-moving HR dynamics while asleep. However, our methods for motion artifact removal, HR curve tracking, and post-processing are unlikely to extend well to scenarios where participants are awake.

## VIII. CONCLUSION

We demonstrated that during sleep wrist-worn accelerometers provide a robust alternative for HR estimation—also in uncontrolled at-home settings where people might share their bed. In accelerometer studies that do not include HR signals, this introduces new opportunities for analysis. For accelerometer studies that were integrated into large longitudinal efforts such as the U.K. Biobank such methods will be of particular value for early disease detection and to enhance our understanding of disease progression and risk factors.

Future work might evaluate the value of Nightbeat for disease modeling (i.e., apply Nightbeat to accelerometer studies of large longitudinal efforts) and extend similar methods to motion-free periods while participants are awake. To extend heart rate estimation techniques to uncontrolled settings while awake, the identification of suitable periods and robust motion artifact removal will likely be crucial. Estimated HR trajectories from such methods might then be used to further understand risk factors of diseases including cardiovascular disease, neurological conditions, or mental disorders.

## IX. LIMITATIONS

Nightbeat-DB, Nightbeat, and our validation has limitations.

First, on both Nightbeat-DB and the AW dataset, we exclude segments when there are no ground-truth HR labels available. This makes the task of HR estimation on the remaining segments seem slightly easier since both ECG and PPG are affected by motion artifacts too. If no HR ground-truth could be derived, this is likely due to strong motion artifacts. This is a general problem of in-the-wild evaluations.

Further, on both datasets, HR labels are relatively low (usually below 90 bpm; see Fig. 5). However, cardiovascular diseases can lead to different HR patterns during the night and the two datasets are thus not fully representative. While Nightbeat does not constraint HR predictions to low ranges, in contrast to the original Bioinsights implementation with filter settings of 0.75–2 Hz, we lack datasets to adequately evaluate approaches for patients with higher HR while asleep.

One drawback of our approach is that it cannot track large HR changes that occur abruptly ( $> 10$  bpm within a few minutes). Given our median smoother during post-processing, such abrupt changes may be tracked by the HR curve and detected in the time domain, but they might be rejected based on the fitted median smooth. While we did not observe such behavior on the two datasets used for validation, it would need evaluation on patients with sleep disorders and cardiovascular diseases to verify Nightbeat's accuracy for such patients.

## REFERENCES

- [1] A. Doherty et al., "Large scale population assessment of physical activity using wrist worn accelerometers: The UK biobank study," *PLoS One*, vol. 12, no. 2, 2017, Art. no. e0169649.
- [2] A.-K. Schalkamp, K. J. Peall, N. A. Harrison, and C. Sandor, "Wearable movement-tracking data identify parkinson's disease years before clinical diagnosis," *Nature Med.*, vol. 29, pp. 2048–2056, 2023.
- [3] K. W. Choi et al., "Assessment of bidirectional relationships between physical activity and depression among adults: A 2-Sample Mendelian Randomization Study," *JAMA Psychiatry*, vol. 76, pp. 399–408, 2019.
- [4] R. Ramakrishnan et al., "Accelerometer measured physical activity and the incidence of cardiovascular disease: Evidence from the UK biobank cohort study," *PLoS Med.*, vol. 18, no. 1, 2021, Art. no. e1003487.
- [5] N. Papadimitriou et al., "Physical activity and risks of breast and colorectal cancer: A Mendelian randomisation analysis," *Nature Commun.*, vol. 11, 2020, Art. no. 597.
- [6] T. Strain et al., "Wearable-device-measured physical activity and future health risk," *Nature Med.*, vol. 26, pp. 1385–1391, 2020.
- [7] M. Marmot and E. Brunner, "Cohort profile: The Whitehall II study," *Int. J. Epidemiol.*, vol. 34, no. 2, pp. 251–256, 2005.
- [8] L. Straker et al., "Cohort profile: The Western Australian pregnancy cohort (Raine) study—generation 2," *Int. J. Epidemiol.*, vol. 46, no. 5, pp. 1384–1385, 2017.
- [9] M. Moebus, M. Hilty, P. Oldrati, L. Barrios, P. A. Consortium, and C. Holz, "Assessing the role of the autonomic nervous system as a driver of sleep quality in patients with multiple sclerosis: Observation study," *JMIR Neurotechnol.*, vol. 3, 2024, Art. no. e48148.
- [10] M. Moebus and C. Holz, "Personalized interpretable prediction of perceived sleep quality: Models with meaningful cardiovascular and behavioral features," *PLoS One*, vol. 19, no. 7, 2024, Art. no. e0305258.
- [11] M. Moebus, C. Holz, and J. Wolfensberger, "Predicting sleep quality via unsupervised learning of cardiac activity," in *Proc. 46th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2024, pp. 15–19.
- [12] R. G. Weaver et al., "Jerks are useful: Extracting pulse rate from wrist-placed accelerometry jerk during sleep in children," *Sleep*, May 2024, Art. no. zsa099, doi: [10.1093/sleep/zsa099](https://doi.org/10.1093/sleep/zsa099).
- [13] J. Zschocke et al., "Reconstruction of pulse wave and respiration from wrist accelerometer during sleep," *IEEE Trans. Biomed. Eng.*, vol. 69, no. 2, pp. 830–839, Feb. 2022.
- [14] J. Zschocke et al., "Detection and analysis of pulse waves during sleep via wrist-worn actigraphy," *PLoS One*, vol. 14, no. 12, 2019, Art. no. e0226843.
- [15] J. Hernandez, D. J. McDuff, and R. W. Picard, "Bioinsights: Extracting personal data from "still" wearable motion sensors," in *Proc. IEEE 12th Int. Conf. Wearable Implantable Body Sensor Netw.*, 2015, pp. 1–6.
- [16] M. Conti, C. Aironi, S. Orcioni, R. Seepold, M. Gaiduk, and N. M. Madrid, "Heart rate detection with accelerometric sensors under the mattress," in *Proc. IEEE 42nd Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2020, pp. 4063–4066.
- [17] V. Bieri, P. Strelci, B. U. Demirel, and C. Holz, "BeliefPPG: Uncertainty-aware heart rate estimation from PPG signals via belief propagation," in *Proc. 39th Conf. Uncertainty Artif. Intell.*, 2023, pp. 173–183.
- [18] M. Zhou and N. Selvaraj, "Heart rate monitoring using sparse spectral curve tracing," in *Proc. IEEE 42nd Annu. Int. Conf. Eng. Med. Biol. Soc.*, 2020, pp. 5347–5352.
- [19] H. Gao, X. Wu, C. Shi, Q. Gao, and J. Geng, "A LSTM-based realtime signal quality assessment for photoplethysmogram and remote photoplethysmogram," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops*, 2021, pp. 3826–3835.
- [20] C. Orphanidou, T. Bonnici, P. Charlton, D. Clifton, D. Vallance, and L. Tarassenko, "Signal-quality indices for the electrocardiogram and photoplethysmogram: Derivation and applications to wireless monitoring," *IEEE J. Biomed. Health Inform.*, vol. 19, no. 3, pp. 832–838, May 2015.
- [21] G. D. Clifford, D. Lopez, Q. Li, and I. Rezek, "Signal quality indices and data fusion for determining acceptability of electrocardiograms collected in noisy ambulatory environments," in *Proc. IEEE Comput. Cardiol.*, 2011, pp. 285–288.
- [22] S. Ancoli-Israel, R. Cole, C. Alessi, M. Chambers, W. Moorcroft, and C. P. Pollak, "The role of actigraphy in the study of sleep and circadian rhythms," *Sleep*, vol. 26, no. 3, pp. 342–392, 2003.
- [23] O. Walch, "Motion and heart rate from a wrist-worn wearable and labeled sleep from polysomnography," *PhysioNet*, vol. 101, 2019, doi: [10.13026/hmhs-py35](https://doi.org/10.13026/hmhs-py35).
- [24] A. L. Goldberger et al., "Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, pp. e215–e220, 2000.
- [25] O. T. Inan et al., "Novel wearable seismocardiography and machine learning algorithms can assess clinical status of heart failure patients," *Circulation, Heart Failure*, vol. 11, no. 1, 2018, Art. no. e004313.
- [26] H. J. Baek, G. S. Chung, K. K. Kim, and K. S. Park, "A smart health monitoring chair for noninvasive measurement of biological signals," *IEEE Trans. Inf. Technol. Biomed.*, vol. 16, no. 1, pp. 150–158, Jan. 2012.
- [27] A. M. Carek and O. T. Inan, "Robust sensing of distal pulse waveforms on a modified weighing scale for ubiquitous pulse transit time measurement," *IEEE Trans. Biomed. Circuits Syst.*, vol. 11, no. 4, pp. 765–772, Aug. 2017.
- [28] C. Hoog Antink et al., "Ballistocardiography can estimate beat-to-beat heart rate accurately at night in patients after vascular intervention," *IEEE J. Biomed. Health Inform.*, vol. 24, no. 8, pp. 2230–2237, Aug. 2020.
- [29] M. Meier and C. Holz, "BMAR: Barometric and motion-based alignment and refinement for offline signal synchronization across devices," *Proc. ACM Interactive, Mobile, Wearable Ubiquitous Technol.*, vol. 7, pp. 1–21, 2023.
- [30] D. Makowski et al., "NeuroKit2: A python toolbox for neurophysiological signal processing," *Behav. Res. Methods*, vol. 53, pp. 1689–1696, 2021.
- [31] J. Pan and W. J. Tompkins, "A real-time QRS detection algorithm," *IEEE Trans. Biomed. Eng.*, vol. BME-32, no. 3, pp. 230–236, Mar. 1985.
- [32] V. Kalidas and L. Tamil, "Real-time QRS detector using stationary wavelet transform for automated ecg analysis," in *Proc. IEEE 17th Int. Conf. Bioinf. Bioeng.*, 2017, pp. 457–461.
- [33] O. Walch, Y. Huang, D. Forger, and C. Goldstein, "Sleep stage prediction with raw acceleration and photoplethysmography heart rate data derived from a consumer wearable device," *Sleep*, vol. 42, no. 12, 2019, Art. no. zsz180.
- [34] B. W. Nelson, C. A. Low, N. Jacobson, P. Areán, J. Torous, and N. B. Allen, "Guidelines for wrist-worn consumer wearable assessment of heart rate in biobehavioral research," *NPJ Digit. Med.*, vol. 3, 2020, Art. no. 90.
- [35] Association for the Advancement of Medical Instrumentation et al., "Cardiac monitors, heart rate meters, and alarms," *American National Standard ANSI/AAMI/EC13:2002: Cardiac Monitors, Heart Rate Meters, and Alarms*, pp. 1–87, 2002.