BCGNet: An AI Model Trained on 600K Hours of Sleep Data for BCG-based Contactless Monitoring

Shigeng Chen[1†], Xuesong Chen[1†], Weijun Huang[2], Fei Lei[3], Chuxuan Shan[1], Zengrui Jin[4], Yunhan Shi[5], Yichen Wang[1], Rui Zhao[6], Xing Xu[7], Dongsheng Lv[6], Yanru Li[5], Brandon Westover[8,9], Shenda Hong[10], Chao Zhang[4], Shankai Yin[2], Chunfeng Liu[7*], Hongliang Yi[2*], Xiangdong Tang[3*], Yue Leng[11*]

[1]Beijing Five Seasons Medical Technology Co., Ltd., Beijing, China.
[2]Department of Otorhinolaryngology Head and Neck Surgery, Shanghai Sixth People's Hospital Affiliated to Shanghai Jiao Tong University School of Medicine, Shanghai, China.
[3]Sleep Medicine Center, West China Hospital, Sichuan University, Chengdu, Sichuan, China.
[4]Department of Electronic Engineering, Tsinghua University, Beijing, China.
[5]Department of Otorhinolaryngology Head and Neck Surgery, Beijing Tongren Hospital, Beijing, China.
[6]Sleep Medicine Center, Inner Mongolia Mental Health Center (The Third Hospital of Inner Mongolia Autonomous Region, Brain Hospital of Inner Mongolia Autonomous Region), Hohhot, Inner Mongolia Autonomous Region, China.
[7]Department of Neurology and Clinical Research Center of Neurological Disease, Second Affiliated Hospital of Soochow University, Suzhou, Jiangsu, China.
[8]Department of Neurology, Harvard Medical School, Boston, Massachusetts, United States.
[9]Department of Neurology, Beth Israel Deaconess Medical Center, Boston, Massachusetts, United States.
[10]National Institute of Health Data Science, Peking University, Beijing, China.
[11]Department of Psychiatry and Behavioral Sciences, University of California, San Francisco, California, United States.

*Corresponding author(s). E-mail(s): liuchunfeng@suda.edu.cn; yihongl@126.com; 2372564613@qq.com; yue.leng@ucsf.edu;
†These authors contributed equally to this work.

Abstract

Sleep profoundly impacts health, yet current gold-standard Polysomnogram (PSG) is constrained by cost, discomfort, and limited scalability for longitudinal monitoring. Ballistocardiogram (BCG) offers a non-invasive and user-friendly alternative but often lacks the precision needed for reliable real-world applications. To address this gap, we propose BCGNet, a two-stage transfer learning model that is first pre-trained on 580,866 hours of PSG and then fine-tuned and validated on 15,081 hours of BCG (total 595,947 hours of recordings). Across multiple validation cohorts, BCGNet achieves strong performance in 4-class sleep staging (F1: $0.710 - 0.817$), Apnea-Hypopnea Index (AHI) estimation (Pearson's $r > 0.95$), and robust quantification of sleep continuity and architecture (ICC and Pearson's r generally > 0.8). Notably, BCGNet maintains strong performance even on short daytime naps and demonstrates excellent generalizability across diverse external datasets. Deployed as a portable, contactless sleep tracking mat, BCGNet represents a major step

towards scalable, user-friendly solutions for longitudinal home sleep monitoring, with important implications for population screening and personalized sleep medicine.

1 Introduction

Sleep profoundly impacts daily health, with poor sleep consistently linked to increased risks of cardiovascular, metabolic, and neuropsychiatric disorders [1–4]. As such, sleep monitoring is essential not only for diagnosing sleep disorders but also for enabling early detection of broader health conditions. Currently, Polysomnogram (PSG) remains the gold standard for evaluating sleep architecture and diagnosing sleep-related disorders [5–7]. However, PSG requires overnight monitoring in a sleep laboratory with multiple sensors attached to the body, which can cause discomfort and introduce first-night effects that may alter sleep patterns [8, 9]. Moreover, standard PSG protocols typically capture only a single night of data, limiting the ability to assess night-to-night variability in habitual sleep patterns [10]. These factors—limited comfort, high cost, and lack of longitudinal capability—make PSG poorly suited for continuous, long-term monitoring in home settings [11, 12].

Ballistocardiogram (BCG) has emerged as a promising non-invasive approach for sleep monitoring, capturing the subtle ballistic forces generated by cardiac ejection, respiratory excursions, and gross body movements during sleep [13, 14]. While BCG offers significant advantages in comfort and feasibility for at-home use [15–19], its reliance on mechanical vibrations makes it highly susceptible to motion artifacts and ambient noise, compromising its reliability in real-world environments [20–23]. This creates a critical unmet need: a technology that combines the diagnostic precision of PSG with the comfort and scalability of home-based sensors.

To address this challenge, we propose a novel transfer learning method that bridges the gap between these two modalities. We hypothesized that a model pre-trained on rich, high-fidelity data from laboratory-grade PSG could be fine-tuned to interpret the sparser, noisier signals of BCG—effectively enabling the model to "learn" complex physiological patterns within the simpler BCG input by transferring knowledge from PSG. To achieve this, we developed BCGNet (Figure 1(c)), a specialized two-stage model based on 580,866 hours of PSG recordings and 15,081 hours of BCG recordings, to robustly translate diagnostic features from the PSG domain into the BCG domain.

Our results, validated across multiple independent clinical cohorts, demonstrate that this approach is highly effective. BCGNet achieves strong performance in 4-class sleep staging (F1: $0.710 - 0.817$), AHI_3 estimation (Pearson's $r > 0.95$, and both Sensitivity and Specificity generally > 0.8), and the quantification of detailed sleep architecture (ICC and Pearson's r generally > 0.8), performing robustly even on short daytime naps. Importantly, the algorithm exhibits robust generalizable and reliability across diverse datasets, marking a significant step towards clinical-grade, contactless sleep assessment. The model is embedded in a portable, non-intrusive sleep tracking mat placed under the user's pillow (Figure 1(a)), offering a practical, scalable, and user-friendly solution for real-world, longitudinal sleep monitoring.

2 Results

2.1 Method Overview

As illustrated in Figure 1(c), our BCGNet harnesses rich pretrained representations through a two-stage transfer learning process from PSG signals. Specifically, during first-stage pre-training, Instantaneous Heart Rate (IHR) and respiration signals, derived respectively from electrocardiogram (ECG) and abdominal (ABD) channels of PSG recordings, are encoded via a ResNet [24] encoder to generate enhanced features. These features are subsequently processed through a RoFormer [25] backbone to capture temporal variations. In the second-stage fine-tuning, BCG-derived IHR and respiration signals are fed into the model layers pretrained on PSG to enable *Channel-Wise Adaptation*. Concurrently, discriminative BCG-specific features, captured via bandpass filtering (marked with asterisks), are adaptively combined with the enhanced features using a *Dynamic Feature Fusion* technique. The fused representation is passed through the RoFormer backbone to generate the final output.

During inference, the model relies solely on BCG signals, including IHR, respiration, and BCG-specific data, for predicting sleep stages and respiratory events. This design allows the model to operate

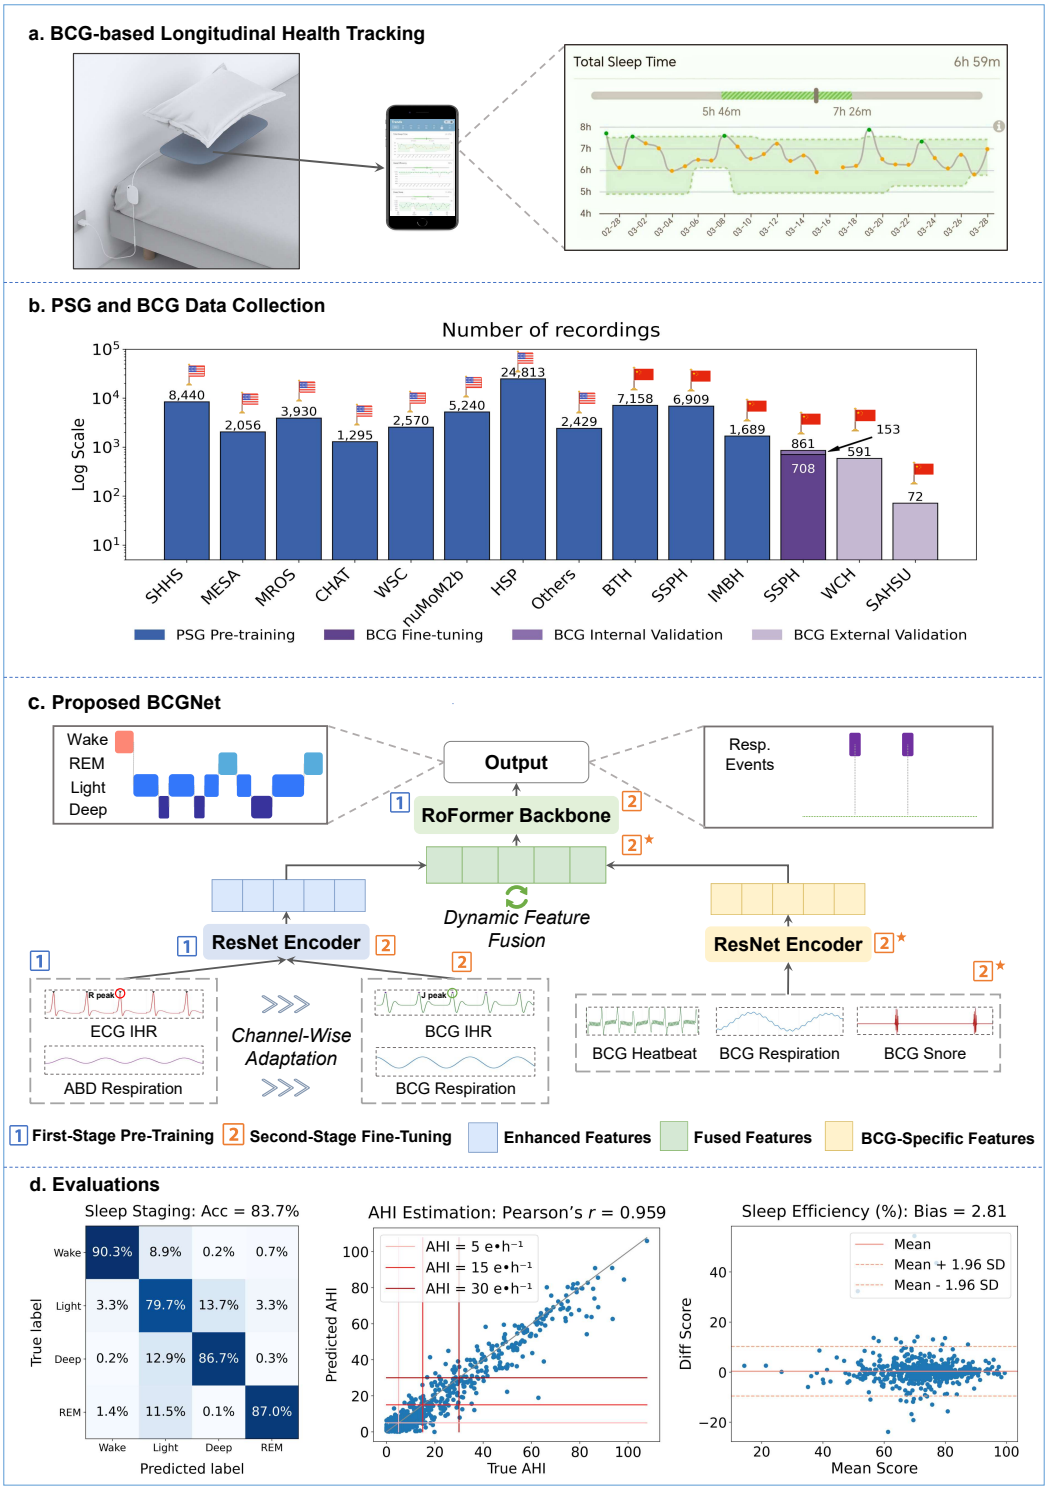

Fig. 1: **Overview of the proposed BCGNet and its validation.** **(a)** Our method of BCG-based longitudinal sleep tracking, where an inner-mattress BCG sensor collects sleep data, and results are visualized on a mobile application for daily monitoring. **(b)** The scale and composition of datasets used in this study. T. PSG cohorts were from the U.S. and China (noted by the flags). The "Others" subsets included CFS, SOF, CCSHS, PATS, and HomePAP from NSRR to diversify and enrich pre-training data. The BCG SSPH was divided such that 708 recordings were utilized for fine-tuning, and 153 recordings were reserved for internal validation. **(c)** The architecture of the proposed BCGNet, which first learns robust representations from ECG and ABD signals, then fine-tunes on BCG signals for final predictions. Details for the training flow in Section 2.1 and training details in Section 4.3 **(d)** Representative evaluation results on the WCH cohort. From left to right: a confusion matrix for 4-class sleep staging, a scatter plot showing the high correlation between predicted and true AHI, and a Bland-Altman plot demonstrating the low bias in sleep efficiency evaluation.

independently of PSG data at deployment, enabling practical, real-world applications where only BCG signals are available.

2.2 Study Population

Table 1: Summary of datasets.

Dataset	Subjects	Recordings	Hours	Female	Age	BMI
PSG						
SHHS	5,795	8,440	75,432	52.8%	64.5 ± 22.3	28.2 ± 10.2
MESA	2,056	2,056	21,745	53.6%	69.4 ± 18.2	28.7 ± 11.2
MROS	2,905	3,930	45,110	0.0%	77.6 ± 11.2	27.1 ± 7.6
CHAT	1,006	1,295	12,884	26.4%	7.0 ± 2.9	19.1 ± 9.8
WSC	1,105	2,570	20,334	46.1%	59.8 ± 17.0	31.7 ± 14.4
nuMoM2b	3,114	5,240	59,795	100.0%	27.0 ± 11.0	27.6 ± 12.8
HSP	18,701	24,813	187,761	43.0%	53.1 ± 33.4	31.6 ± 15.3
BTH	6,327	7,158	58,639	22.7%	34.5 ± 17.8	24.8 ± 5.4
SSPH	7,009	6,909	61,504	20.9%	40.5 ± 25.8	26.3 ± 13.2
IMBH	1,689	1,689	12,982	63.6%	52.0 ± 30.4	24.1 ± 7.7
BCG						
SSPH	861	861	8,708	27.6%	42.2 ± 28.0	25.9 ± 8.3
WCH	591	591	5,872	25.8%	42.8 ± 30.8	–
SAHSU	72	72	727	18.1%	44.5 ± 32.5	26.6 ± 8.1

Subjects = number of participants; Hours = total recording duration in hours; Age and BMI are presented as Mean ± 2SD.

Table 1 summarizes the demographic and anthropometric characteristics of the main cohorts used in this study. A key strength of our study is the unprecedented scale and diversity of the data collected.

In total, we used 595,947 hours of sleep recordings to develop BCGNet. For the crucial PSG pre-training stage, we aggregated a massive and diverse dataset comprising 580,866 hours of overnight recordings. This pre-training dataset was intentionally curated for diversity, sourced from over ten distinct cohorts. These included large-scale, publicly available U.S. datasets from the Human Sleep Project (HSP) and the National Sleep Research Resource (NSRR). This pre-training dataset was further augmented by cohorts from hospitals in China, including Beijing Tongren Hospital (BTH), Shanghai Sixth People's Hospital (SSPH), and Inner Mongolia Brain Hospital (IMBH), a cohort specializing in psychiatric-related conditions such as depression and insomnia. This inclusion ensures the model's exposure to a vast spectrum of sleep phenotypes and clinical conditions, thereby underscoring the remarkable diversity of our dataset.

For the subsequent fine-tuning and validation stages, we collected 15,081-hours BCG recordings from three independent hospitals in China: the SSPH cohort (N=861; 8,482 hours), which was partitioned into a fine-tuning set (N=708; 6,960 hours) and an internal validation set (N=153, 1,522 hours); and two external validation cohorts from the West China Hospital (WCH; N=591; 5,872 hours) and The Second Affiliated Hospital of Soochow University (SAHSU; N=72; 727 hours). By splitting data by subject across fine-tuning and validation sets, we ensure that BCGNet's final performance is robust and generalizable in both internal and external validation. The sheer scale and heterogeneity of these datasets are foundational to BCGNet's high accuracy and reliability.

2.3 BCGNet is comparable to PSG in sleep staging

Table 2: **Sleep staging performance across multi-center cohorts**

Cohorts	Cohen's κ	Accuracy	F1
SSPH	0.766 (0.750,0.781)	0.844 (0.834,0.854)	0.817 (0.804,0.831)
WCH	0.736 (0.728,0.743)	0.837 (0.832,0.842)	0.710 (0.700,0.718)
SAHSU	0.697 (0.666,0.725)	0.797 (0.778,0.815)	0.781 (0.759,0.800)

Our BCGNet model demonstrated robust performance in 4-class sleep staging across multiple independent cohorts, as summarized in Table 2 and visualized by the confusion matrices in Figure 2. The internal validation on the SSPH cohort yielded the highest agreement, achieving a Cohen's κ of 0.766 (95% CI: 0.750, 0.781) and an F1 of 0.817 (95% CI: 0.804, 0.831). While a slight reduction in performance was observed in the external validation cohorts, WCH and SAHSU, the model maintained strong and clinically acceptable performance. Specifically, BCGNet achieved a Cohen's κ of 0.736 (95% CI: 0.728, 0.743) and F1 of 0.710 (95% CI: 0.700, 0.718) in the WCH, and a Cohen's κ of 0.697 (95% CI: 0.666, 0.725) and F1 of 0.781 (95% CI: 0.759, 0.800) in the SAHSU. The 3-class and 5-class sleep staging results can be found in Supplementary Figure 1. This multi-center validation underscores the generalizability and robustness of BCGNet for reliable sleep staging across multi-center cohorts.

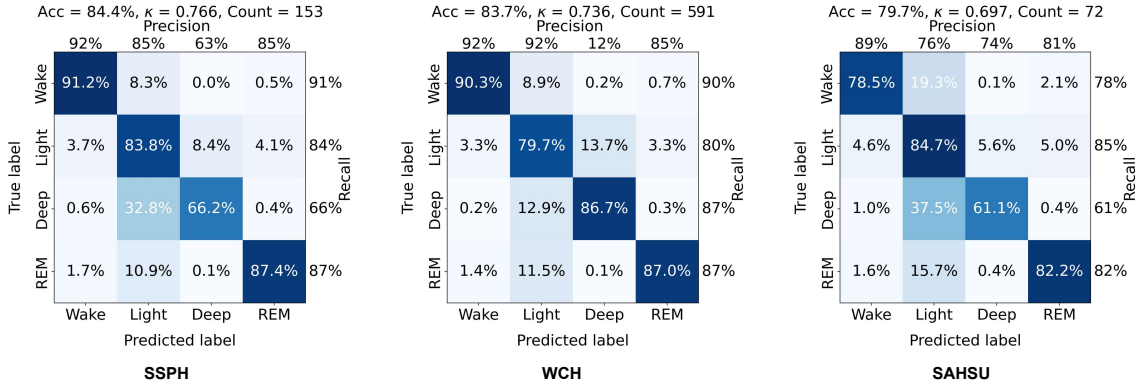

Fig. 2: **Performance of BCGNet for 4-class sleep staging across multi-center cohorts.**

2.4 BCGNet robustly estimates AHI3 across multi-center cohorts

For continuous AHI_3 estimation, the scatter plots in the top row of Figure 3 reveal a remarkably strong linear correlation between predicted and true AHI_3 values across all cohorts, with Pearson's r consistently exceeding 0.95 (SSPH: $r = 0.963$; WCH: $r = 0.959$; SAHSU: $r = 0.951$). The Bland-Altman plots in the bottom row of Figure 3 further illustrate the level of agreement, showing minimal mean biases (SSPH: -1.63 events/h; WCH: -2.55 events/h; SAHSU: +0.725 events/h). The 95% limits of agreement (LoA) remained relatively tight for SSPH and WCH, indicating strong consistency between predicted and reference values. While SAHSU exhibited slightly wider limits, the overall low bias and strong correlation across all cohorts highlight the model's robust calibration for AHI_3 prediction in diverse clinical settings.

The BCGNet also achieved high accuracy in sleep apnea severity classification (Healthy, Mild, Moderate, Severe), as shown in Supplementary Figure 2. In the SSPH cohort, the highest values along the confusion-matrix diagonal, such as 94.3 % for Healthy and 82.5 % for Severe, demonstrate BCGNet's strong classification performance across all categories. Similar performance patterns were observed in the WCH and SAHSU cohorts. These results highlight BCGNet's capability in identifying clinically significant sleep apnea across diverse populations.

Furthermore, the bar charts in Supplementary Figure 3 demonstrate the model's high diagnostic utility across standard AHI thresholds (5, 15, and 30 events/h). BCGNet consistently achieved high sensitivity and specificity for detecting mild, moderate, and severe sleep apnea levels across all evaluated cohorts. This robust and consistent performance at clinically meaningful cutoffs reinforces the model's potential for accurate and reliable automated sleep apnea assessment in real-world applications.

2.5 BCGNet accurately estimates overnight sleep continuity and architecture

Beyond sleep staging and AHI_3 estimation, BCGNet also demonstrated strong performance in quantifying key sleep continuity and architecture parameters from full-night recordings. These included Total

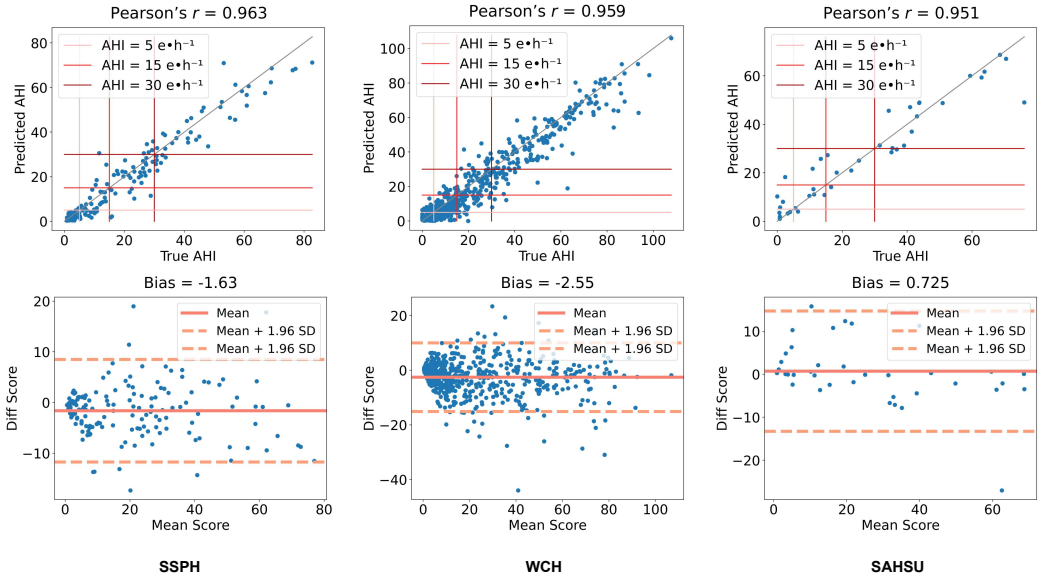

Fig. 3: **Comprehensive performance of BCGNet for AHI estimation across multi-center cohorts. Top Row**: Scatter plots showing predicted AHI_3 versus true AHI_3 values. **Bottom Row**: Bland-Altman plots depicting the mean bias and 95% limits of agreement between predicted and true AHI_3 values.

Sleep Time (TST), Sleep Onset Latency (SOL), Wake After Sleep Onset (WASO), Sleep Efficiency (SE), and REM Latency (REML).

As shown in Figure 4, the model consistently achieved high Intraclass Correlation Coefficient (ICC) and Pearson's correlation coefficients (r) across all evaluated sleep parameters and cohorts. In the internal SSPH cohort, both ICC and Pearson's r values for TST, SOL, WASO, SE, and REML generally exceeded 0.8, with TST and SE showing particularly strong agreement (both ICC and $r > 0.97$). While a slight decrease in correlation was observed in the external WCH and SAHSU cohorts, the performance remained robust, with most ICC and r values well above 0.8, indicating strong reliability and linear association between predicted and ground-truth values.

The Bland-Altman plots in Supplementary Figure 4 provide a detailed visualization of the agreement between BCGNet predictions and expert annotations for each sleep parameter across the SSPH, WCH, and SAHSU cohorts. Across all parameters and cohorts, the plots consistently reveal minimal mean biases, as indicated by the solid red line close to zero, suggesting that model predictions closely align with ground-truth values. The 95% LoA, as indicated by dashed red lines, demonstrates that the majority of the data points fall within acceptable bounds, reflecting a high degree of precision and consistency. As expected, parameters with greater inherent variability, such as SOL and WASO, exhibited slightly wider LoA. In contrast, parameters like TST and SE showed tight clustering around the mean difference line, underscoring strong agreement.

2.6 BCGNet supports precise monitoring of daytime naps

Beyond the comprehensive assessment of overnight sleep, we further evaluated the performance of BCGNet on daytime naps. For this evaluation, we utilized data from a separate WCH cohort specifically dedicated to multiple sleep latency test (MSLT) studies. Due to the limited presence or absence of deeper sleep stages (N3 and REM) in typical short naps, we focused on three primary sleep continuity parameters: TST, SOL, and SE.

As presented in Table 5(a) and further illustrated by the 2-class sleep staging confusion matrix in Figure 5(b), our model demonstrated reliable performance for short naps. For the selected metrics, BCGNet achieved an ICC of 0.717 (95% CI: 0.663, 0.765) and a Pearson's r of 0.725 (95% CI: 0.673, 0.773) for TST. SOL showed an ICC of 0.638 (95% CI: 0.578, 0.693) and a Pearson's r of 0.653 (95% CI: 0.590, 0.712). Similarly, SE yielded an ICC of 0.669 (95% CI: 0.620, 0.709) and a Pearson's r of 0.672 (95% CI: 0.623, 0.714). These results indicate a good to substantial agreement and correlation, confirming the model's capability to quantify these essential parameters even in short sleep episodes.

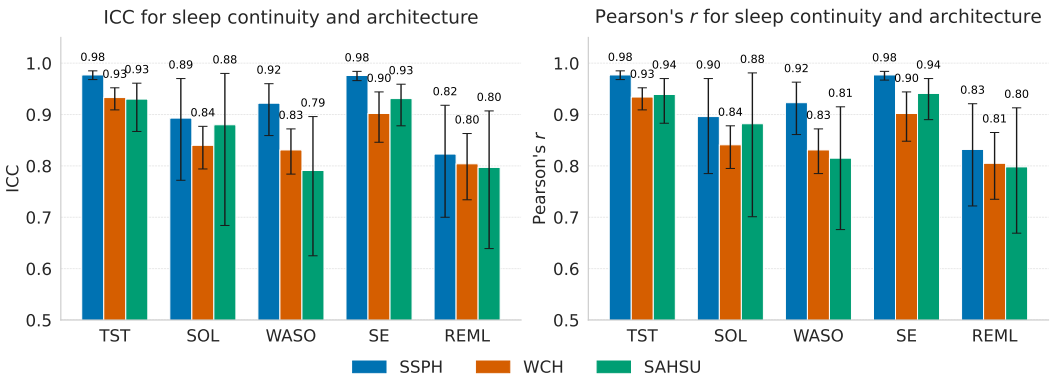

Fig. 4: **Performance of BCGNet in estimating overnight sleep continuity and architecture parameters.** This figure presents the ICC and Pearson's correlation coefficient (r) for five key sleep parameters: TST, SOL, WASO, SE, and REML. Results are grouped by parameter and displayed for three cohorts: SSPH, WCH, and SAHSU. Error bars represent the 95% confidence intervals.

Furthermore, the Sleep-Wake staging for short naps (Figure 5(b)) highlights the model's robust ability to differentiate between wakefulness and sleep states. With an overall accuracy of 81.3% and a Cohen's κ of 0.624, the model achieved high precision (83% for Wake, 80% for Sleep) and recall (76% for Wake, 86% for Sleep). This strong binary classification performance is crucial for accurately determining the presence and duration of daytime naps.

In summary, BCGNet exhibits exceptional versatility and accuracy, extending its robust performance from comprehensive overnight sleep analysis to the nuanced assessment of daytime naps, making it a powerful tool for continuous, daily sleep monitoring across various settings.

Metrics	ICC	Pearson's r
TST	0.717 (0.663, 0.765)	0.725 (0.673, 0.773)
SOL	0.638 (0.578, 0.693)	0.653 (0.590, 0.712)
SE	0.669 (0.620, 0.709)	0.672 (0.623, 0.714)

a. Continuity performance.

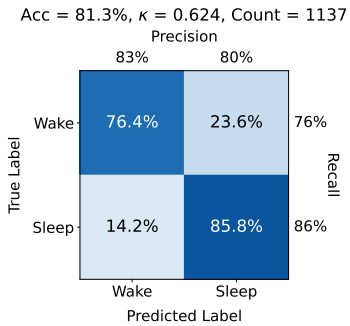

b. Sleep–Wake staging.

Fig. 5: **Evaluation of BCGNet performance on daytime naps. (a)** For short naps, Total Sleep Time (TST), Sleep Onset Latency (SOL), and Sleep Efficiency (SE) are included as sleep continuity metrics. **(b)** Sleep–Wake staging evaluation is performed on the short-snap cohort.

2.7 Ablation study

Ablation experiments were conducted to elucidate the contribution of each key component within the proposed BCGNet to its overall performance. Specifically, we compared three distinct training schemes: (i) BCG-based, a model trained purely on BCG signals without PSG pre-training; (ii) PSG-based, representing a zero-shot condition where the PSG pre-trained model directly inferred on BCG data without fine-tuning; and (iii) Proposed BCGNet, our complete two-stage transfer learning model.

As presented in Table 3, the BCG-based model achieved a Cohen's κ of 0.629 and an F1 of 0.717 for sleep staging. For AHI_3 estimation, it yielded a Pearson's r of 0.951 and an F1 of 0.713. The PSG-based model, while leveraging PSG-learned representations but not fine-tuning on BCG, showed reduced performance with Cohen's κ of 0.742 and F1 of 0.795 for sleep staging, and Pearson's r of 0.949 and F1 of 0.666 for AHI_3 estimation. Notably, our proposed BCGNet generally demonstrated

Table 3: **Results of ablation study on sleep staging and AHI_3 estimation.**

Method	**Sleep staging**		AHI_3 **estimation**	
	Cohen's κ	**F1**	**Pearson's r**	**F1**
BCG-based	0.629 (0.607,0.648)	0.717 (0.702,0.732)	0.951 (0.932, 0.966)	**0.713** **(0.640,0.780)**
PSG-based (*w/o* fine-tuning)	0.742 (0.724,0.760)	0.795 (0.780,0.809)	0.949 (0.925, 0.966)	0.666 (0.590,0.732)
BCGNet (proposed)	**0.766** **(0.750,0.781)**	**0.817** **(0.804,0.831)**	**0.963** **(0.946, 0.975)**	0.711 0.631,0.775)

"*w/o*" stands for "without". The results are based on the internal cohort SSPH. The bolded values indicate the highest performance for each task.

superior performance, achieving Cohen's κ of 0.766 and F1 of 0.817 for sleep staging, and Pearson's r (0.963) and F1 (0.711) for AHI_3 estimation. Overall, these findings highlight that the two-stage training approach substantially improves the model performance by leveraging the transferred knowledge.

3 Discussion

In this study, we developed and validated BCGNet, a state-of-the-art transfer learning model enables comprehensive sleep assessment across multiple settings . Our findings demonstrate that BCGNet achieves high accuracy and robustness across multi-center cohorts for sleep staging, AHI estimation, assessment of key sleep continuity and architecture parameters, and analysis of short daytime naps. These results represent a significant step towards scalable, longitudinal, and contactless sleep monitoring in naturalistic settings.

Table 4: Related work in the literature.

Source (Year)	**Sensor**	**Sleep staging**	**AHI estimation**
Korkalainen et al. (2020)[26]	PPG (full signals)	3-cls: Acc=0.80, κ=0.65 4-cls: Acc=0.69 κ=0.54 5-cls: Acc=0.64, κ=0.51	–
Sun et al. (2020) [27]	ECG, ABD (HR & Resp)	5-cls: κ=0.53, F1=0.59	–
Sridhar et al. (2020) [28]	ECG (IHR)	4-cls: Acc=0.77, κ=0.66	–
Radha et al. (2021) [29]	PPG (HRV)	4-cls: Acc=0.76, κ=0.65	–
Li et al. (2024) [30]	BCG (HR & Resp)	5-cls: Acc=0.70, κ=0.59	–
Liu et al. (2024) [31]	ECG (full signals)	–	cufoff at 5: Acc=0.81 cufoff at 15: Acc=0.91 cutoff at 30: Acc=0.93
Zhang et al. (2025) [32]	ECG, SpO_2 (full signals)	–	cufoff at 5: Acc=0.90 cufoff at 15: Acc=0.88 cutoff at 30: Acc=0.92
He et al. (2025) [33]	Radio (wave)	4-cls: Acc=0.80, κ=0.68	r=0.82 cutoff at 5: Sens=0.74, Spec=0.89 cutoff at 15: Sens=0.73, Spec=0.95
Ours	BCG (IHR & Resp)	3-cls: Acc=0.91, κ=0.83, F1=0.89 4-cls: Acc=0.83, κ=0.73, F1=0.77 5-cls: Acc=0.76, κ=0.67, F1=0.70	r=0.95, 4-cls: Acc=0.73, F1=0.75 cutoff at 5: Sens=0.90, Spec=0.86, Acc=0.89 15: Sens=0.90, Spec=0.90, Acc=0.91 30: Sens=0.85, Spec=0.97, Acc=0.93

"cls" stands for classification. For example, in sleep staging, "4-cls" refers to four-class staging; In AHI estimation, "4-cls" represents the four categories: healthy, mild, moderate, and severe. "full signals" refers to using the raw, continuous signals as the input to the model. "HR" and "Resp" refer to Heart Rate and Respiration, respectively. "HRV" and "IHR" are abbreviations for Heart Rate Variability and Instantaneous Heart Rate. For specific details on how these features are processed and extracted, please refer to the original papers.

Compared to recent published studies (summarized in Table 4), our approach differs fundamentally in scale (595,947-hour recordings in total) and methodology, leading to superior performance and robustness. The proposed BCGNet not only outperforms previous BCG based methods but also achieves performance comparable to or exceeding that of approaches based on PPG, radio, or ECG. A detailed comparison reveals several key advantages:

First, the scale and diversity of our training data are unparalleled. Previous studies on BCG-based sleep staging or apnea detection typically rely on smaller, single-center datasets, often numbering in the dozens or over one hundred subjects [34, 35]. Such models risk overfitting to local population characteristics and often fail to generalize beyond their training data [36, 37]. In contrast, our diverse and large training data exposed the BCGNet to a vast spectrum of sleep phenotypes, demographic variations (age, sex, BMI, ethnicity), and technical variabilities, providing it with a deep and generalizable understanding of sleep patterns that is foundational to its robustness.

Second, we introduce a paradigm shift from direct training to transfer learning. Typically, transfer learning from PSG to BCG is challenging because BCG recordings exhibit a substantially lower signal-to-noise ratio (SNR) than PSG, which reduces the fidelity of directly transferred representations. However, our transfer learning model overcomes this limitation through a two-stage procedure. With *Channel-Wise Adaptation*, the model first learns robust, time-aware features from dual-channel PSG data (ECG and ABD). Subsequently, *Dynamic Feature Fusion* adaptively integrates these aligned high-level PSG-derived features with the raw continuous BCG-specific features from BCG, ensuring that the transferred information is both meaningful and usable despite the poorer SNR of BCG. As a result, our approach outperforms most existing models that are trained directly and exclusively on a single modality [30, 38, 39].

More importantly, our validation is more comprehensive. We evaluate the model using metrics for both sleep staging and AHI estimation (see Table 4), as well as sleep architecture parameters (see Figure 4). In contrast, previous studies have typically focused on only a subset of these measures [28, 29, 31, 32, 40–45]. Our thorough assessment demonstrates that BCGNet provides a broader, more complete, and more accurate solution for contactless sleep monitoring. Furthermore, validation on daytime naps broadens its applicability and potential for capturing a holistic picture of an individual's 24-hour sleep-wake patterns, which are increasingly recognized as crucial for health and well-being [46, 47].

We acknowledge several limitations. First, the validation cohorts were composed of Chinese individuals referred for sleep studies. Future work aims to validate BCGNet on more ethnically and geographically diverse populations. Second, the assessment of short naps requires the user to be lying on the bed where the sleep monitoring mat is placed, which limits its application for tracking naps in other locations (e.g., on a sofa). Addressing this will require the development of portable or non-bed-based sensors capable of maintaining signal fidelity outside the traditional sleep environment. Additionally, although our model achieved high performance on the external cohorts, we observed a slight reduction in sleep staging when compared to its performance on the internal cohort. We speculate that this discrepancy may arise from the inherent variability in sleep stage scoring across different centers [48]. In our future work, we plan to incorporate analysis based on annotations from multiple experts to better understand and mitigate the impact of such inconsistencies.

In conclusion, our BCGNet constitutes a notable methodological advance in contactless sleep monitoring. By effectively transferring knowledge from PSG to BCG signals, it achieves superior accuracy and coverage using entirely non-intrusive method. Our findings lay the groundwork for scalable, clinically validated tools that enable comfortable, continuous assessment of sleep health.

4 Methods

4.1 Datasets

We collected publicly available PSG datasets from multiple large-scale sleep studies to ensure diversity and robustness in the first-stage pre-training:

- **Human Sleep Project (HSP)**[1] comprises 24,859 PSG recordings from 18,701 patients evaluated at the Massachusetts General Hospital Sleep Division, with signals acquired per AASM standards (such as EEG, ECG, EOG, EMG, and airflow).

[1] https://bdsp.io/content/hsp/2.0

- **National Sleep Research Resource (NSRR)**[2] provides a comprehensive repository of polysomnographic (PSG) data, demographic records, and clinical annotations from multicenter cohorts. We used the following main PSG subdatasets: Sleep Heart Health Study (SHHS), Multi-Ethnic Study of Atherosclerosis (MESA), Osteoporotic Fractures in Men Study (MROS), Childhood Adenotonsillectomy Trial (CHAT), Wisconsin Sleep Cohort (WSC), Nulliparous Pregnancy Outcomes Study: Monitoring Mothers-to-Be (nuMoM2b), Cleveland Family Study (CFS), Study of Osteoporotic Fractures (SOF), Cleveland Children's Sleep and Health Study (CCSHS), Pediatric Adenotonsillectomy Trial of Snoring (PATS), and Home Positive Airway Pressure (HomePAP),
- **Shanghai Sixth People's Hospital (SSPH)** comprises 7,009 overnight PSG recordings primarily from patients with Obstructive Sleep Apnea (OSA), including comprehensive annotations for sleep stages and respiratory events. Additionally, 861 BCG recordings from this cohort were utilized in this study.
- **Beijing Tongren Hospital (BTH)** consists of 6,327 overnight PSG recordings, mainly from the Otorhinolaryngology Department, focusing on patients with suspected OSA.
- **Inner Mongolia Brain Hospital (IMBH)** comprises 1,689 PSG recordings collected at its sleep center from patients diagnosed with psychiatric disorders and obstructive sleep apnea.
- **The Second Affiliated Hospital of Soochow University (SAHSU)** comprises 72 BCG recordings from its Neurology Department. Due to missing AHI labels, 40 of these recordings were ultimately used for sleep apnea validation.
- **West China Hospital (WCH)** includes 591 BCG recordings from its Sleep Medicine Center.

4.2 Signal extraction

Signal Quality Assessment and Selection: To ensure the reliability of PSG and BCG recordings used in model training, we implement a signal quality assessment process based on the energy distribution within respiratory and heartbeat frequency bands. Specifically, for each time segment, we compute the proportion of total signal energy contained within the respiration and heartbeat bands, denoted as the confidence score C:

$$C = \frac{E_{\text{respiration}} + E_{\text{heartbeat}}}{E_{\text{total}}}, \quad C \in [0, 1] \tag{1}$$

where $E_{\text{respiration}}$ and $E_{\text{heartbeat}}$ are the energy components in the respiration and heartbeat frequency ranges, respectively, and E_{total} represents the total signal energy in the segment.

For each recording, an average confidence score $\bar{C}$ is computed over all time segments. A recording is considered valid if $\bar{C} > 0.5$. Instead of discarding low-quality segments within a recording, we apply this threshold at the recording level, ensuring that only recordings with sufficiently high overall signal quality are included in the dataset. By maintaining a comparable quality control strategy across both PSG and BCG modalities, we improve the robustness of the model and prevent the introduction of low-quality physiological signals that may negatively impact training and evaluation.

Then, we process the original signal through breath signal extraction, heartbeat signal extraction, and body movement removal processes.

Respiratory signal extraction: The respiratory signal is derived from both the BCG and ABD signals in the PSG. First, a third-order Bessel band-pass filter [49] (0.1–0.35 Hz) is applied to the original BCG and ABD signals. The filtered signals are resampled to 4 Hz. For the BCG-derived breath signal, temporal integration is performed to address phase discrepancies between the BCG (representing breathing flux) and ABD (reflecting air volume). Integration is omitted for the PSG-derived signal. Finally, both signals are standardized to z-scores to eliminate magnitude differences inherent to sensor principles.

Heartbeat signal extraction: Heartbeat intervals are represented as inter-beat intervals (IBIs). For BCG signals, J peaks define JJ intervals (JJIs), while R peaks in ECG signals represent RR intervals (RRIs). The BCG-derived heartbeats are extracted using the approach proposed by [50]. For ECG, R peaks are detected via NeuroKit2 [51]. Outliers are replaced using HRVAnalysis [52], which removes physiologically implausible intervals via linear interpolation. Removed intervals at sequence extremities are truncated. The IBIs are linearly interpolated at 4 Hz to form a continuous signal. The resulting JJIs and RRIs, expressed in milliseconds, are then converted to IHRs in beats per minute (BPM) via $1000 \times 60/\text{JJI}$, and the IHRs are used as model inputs.

[2] https://sleepdata.org

Body movement detection: Body movements induce rapid fluctuations in BCG and EEG signals. Given the 500 Hz sampling rate of these signals, a 2-second window is used to capture transient variations. The peak-to-peak amplitude within each window serves as the movement feature. A dynamic 30-second baseline (divided into 15 consecutive 2-second windows) is established to compute the mean (μ) and standard deviation (σ) of the feature sequence. Significant deviations are identified using a threshold of $\mu + 5\sigma$, which minimizes false positives by excluding physiological fluctuations. This threshold aligns with statistical norms for Gaussian-distributed data, ensuring robust detection of movement-induced artifacts.

4.3 Model Development

We propose a two-stage transfer learning model to address sleep staging and AHI_3 estimation, focusing on two different data sources: PSG and BCG.

Stage 1 (Pre-training on PSG).

In the first stage, we train our model using two channels from PSG recordings: Heartbeat (ECG-based RRI) and Respiration (ABD belt).

Each channel is fed into a dedicated 1D ResNet-like [24] feature extractor, which consists of: 1) A convolutional layer (kernel size = 7, padding = 3), followed by batch normalization and ReLU activation. 2) A max pooling layer (kernel size = 3, padding = 1). and 3) Four sequential residual blocks (each containing two `ResBlock1D` units).

The extracted feature representations from the two channels are summed element-wise to form a combined feature. We prepend a learnable `CLS` token to this tensor and feed the resulting sequence into a RoFormer [25] with an embedding size of 512. A final fully connected layer produces the logits for sleep staging or AHI_3 estimation. This stage learns robust representations of the heartbeat and respiration signals from the PSG domain.

Stage 2 (Fine-tuning on BCG).

In this stage, we fine-tune the pre-trained model on BCG. The network now ingests two *types* of input channels:

- **Transferred channels** (from the PSG model):
  - **BCG IHR** – weights transferred from the PSG heartbeat branch.
  - **BCG Respiration** – weights transferred from the PSG respiration branch.
- **BCG-specific channels** (each with its own 1D ResNet extractor):
  - **BCG Respiration** (0.1–1 Hz, sampling rate = 2 Hz)
  - **BCG Heartbeat** (0.5–25 Hz, sampling rate = 50 Hz)
  - **BCG Snore** (30–200 Hz, sampling rate = 400 Hz)

The two transferred channels leverage the pretrained PSG weights directly. Each BCG-specific channel is processed by a separate 1D ResNet feature extractor—architecturally identical to the Stage 1 extractors but adapted to its signal's token length. A learnable gating mechanism then adaptively weighs the three BCG-specific feature streams, which are fused with the two transferred channels via element-wise addition to form a single representation. After prepending the `CLS` token, this fused feature is passed through the same RoFormer encoder and final fully-connected layer, allowing the model to exploit both transferred PSG knowledge and the additional discriminative information in the BCG-specific BCG signals.

We evaluate the proposed BCGNet on two tasks: *sleep staging* and AHI_3 *estimation*. Table 5 summarizes the key training hyperparameters for both tasks. In terms of sleep staging setup, each input segment is 30 s in duration (`token_length` = 30 s). The entire night recording is split into consecutive 30 s segments for training and evaluation. For AHI_3 estimation, we segment the data at 1 s intervals (`token_length` = 1 s) and train on 30 min segments (`input_length` = 30 min). We use the Adam optimizer with an initial learning rate of 1×10^{-4} for both tasks, reduced to 1×10^{-5} during fine-tuning. Weighted Cross Entropy is adopted for sleep staging and Focal Loss [53] for AHI estimation, where apnea and hypopnea are grouped into one class, and normal breathing events as another class. For sleep staging, the dataset was partitioned into training, validation, and test sets in an 8:1:1 ratio. We selected balanced validation and test subsets, each containing 4,000 samples, considering the class imbalance for the AHI estimation task. The codes are available in section 4.4.

Table 5: Key training hyperparameters for both tasks.

Setting	**Sleep Staging**	AHI_3 **Estimation**
Optimizer	Adam	Adam
Loss function	Weighted Cross Entropy Loss	Focal Loss
Batch size	32	32
Token length	30 s	1 s
Input length	Whole night	30 min

4.4 Evaluations

Sleep staging metrics.

For sleep staging, we use Cohen's kappa (κ), Accuracy (Acc), and F1, as overall performance indicators, while Precision and Recall are reported for each sleep stage and presented in a confusion matrix.

- κ accounts for chance agreement between predictions and ground truth:

$$\kappa = \frac{P_o - P_e}{1 - P_e} \tag{2}$$

  where P_o is the observed agreement and P_e is the expected agreement under random chance.
- **Acc** measures the overall proportion of correctly classified sleep stages:

$$\text{Acc} = \frac{TP + TN}{TP + TN + FP + FN}, \tag{3}$$

  where TP, TN, FP, and FN denote the counts of true positives, true negatives, false positives, and false negatives, respectively.
- **Precision** for each sleep stage quantifies the fraction of epochs predicted as that stage that are truly of that stage:

$$\text{Precision} = \frac{TP}{TP + FP}. \tag{4}$$

- **Recall** for each sleep stage indicates the fraction of actual epochs of that stage that are correctly identified:

$$\text{Recall} = \frac{TP}{TP + FN}. \tag{5}$$

- **F1**: For each sleep stage i, the F1 is the harmonic mean of its precision and recall:

$$\text{F1}_i = 2 \times \frac{\text{Precision}_i \times \text{Recall}_i}{\text{Precision}_i + \text{Recall}_i}. \tag{6}$$

  In this paper, we used the Macro-F1 as default, which computes the average of the F1-scores for each stage. This approach treats all classes equally, regardless of their prevalence in the dataset:

$$\text{Macro-F1} = \frac{1}{K} \sum_{i=1}^{K} \text{F1}_i, \tag{7}$$

  where K is the number of sleep stages.

AHI estimation metrics.

We report Sensitivity (Sens) and Specificity (Spec) across three AHI severity thresholds: mild–5, moderate–15, and severe–30 sleep apnea, based on standard clinical cutoffs. For global AHI estimation, we use the Pearson correlation coefficient (r) as the primary evaluation metric, complemented by Scatter plots and Bland–Altman plots for visual assessment. At the 4-class classification level, we evaluate performance using macro F1 score and confusion matrices to capture class-wise precision and recall balance.

- **Sens** measures the proportion of correctly detected positive cases:

$$\text{Sens} = \frac{TP}{TP + FN} \tag{8}$$

- **Spec** measures the proportion of correctly detected negative cases:

$$\text{Spec} = \frac{TN}{TN + FP} \tag{9}$$

- r quantifies the linear correlation between the predicted AHI values and the ground truth AHI values:

$$r = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}} \tag{10}$$

Sleep continuity and architecture metrics.

To quantify the accuracy and consistency of our model in estimating macro-architecture parameters, we evaluate the following five metrics: Total Sleep Time (TST), Sleep Onset Latency (SOL), Wake After Sleep Onset (WASO), Sleep Efficiency (SE), and REM Latency (REML). We assess each of the five sleep-quality measures with Pearson r and ICC to provide a complete picture of linear agreement and measurement reliability.

- **Intraclass Correlation Coefficient (ICC)**:

$$\text{ICC} = \frac{MS_B - MS_W}{MS_B + (k-1)\,MS_W + k\,\frac{MS_R - MS_W}{n}}, \tag{11}$$

  where MS_B, MS_W, and MS_R are the between-subjects, within-subjects, and residual mean squares, respectively, k is the number of measurements per subject (here $k = 2$, the true and predicted values), and n is the number of subjects.

Data availability

The datasets from NSRR are obtained via https://sleepdata.org, and the HSP datasets are available at https://bdsp.io/content/hsp/2.0. The Chinese cohorts are available from the authors upon reasonable request.

Code availability

The codes are open source and available at https://github.com/FiveSeasonsMedicine/BCGNet.

References

[1] Lim, D. C. et al. The need to promote sleep health in public health agendas across the globe. *Lancet Public Health* **8**, e820–e826 (2023).

[2] Irwin, M. R. Why sleep is important for health: a psychoneuroimmunology perspective. *Annu. Rev. Psychol.* **66**, 143–172 (2015).

[3] Mukherjee, S. et al. An official american thoracic society statement: the importance of healthy sleep. recommendations and future priorities. *Am. J. Respir. Crit. Care Med.* **191**, 1450–1458 (2015).

[4] Leng, Y., Musiek, E. S., Hu, K., Cappuccio, F. P. & Yaffe, K. Association between circadian rhythms and neurodegenerative diseases. *Lancet Neurol.* **18**, 307–318 (2019).

[5] Rundo, J. V. & Downey III, R. Polysomnography. *Handb. Clin. Neurol.* **160**, 381–392 (2019).

[6] Bloch, K. E. Polysomnography: a systematic review. *Technol. Health Care* **5**, 285–305 (1997).

[7] Boulos, M. I. et al. Normal polysomnography parameters in healthy adults: a systematic review and meta-analysis. *Lancet Respir. Med.* **7**, 533–543 (2019).

[8] Portier, F. et al. Evaluation of home versus laboratory polysomnography in the diagnosis of sleep apnea syndrome. *Am. J. Respir. Crit. Care Med.* **162**, 814–818 (2000).

[9] Miettinen, T. et al. Home polysomnography reveals a first-night effect in patients with low sleep bruxism activity. *J. Clin. Sleep Med.* **14**, 1377–1386 (2018).

[10] Tschopp, S., Wimmer, W., Caversaccio, M., Borner, U. & Tschopp, K. Night-to-night variability in obstructive sleep apnea using peripheral arterial tonometry: a case for multiple night testing. *J. Clin. Sleep Med.* **17**, 1751–1758 (2021).

[11] Caples, S. M., Anderson, W. M., Calero, K., Howell, M. & Hashmi, S. D. Use of polysomnography and home sleep apnea tests for the longitudinal management of obstructive sleep apnea in adults: an american academy of sleep medicine clinical guidance statement. *J. Clin. Sleep Med.* **17**, 1287–1293 (2021).

[12] Birrer, V., Elgendi, M., Lambercy, O. & Menon, C. Evaluating reliability in wearable devices for sleep staging. *NPJ Digit. Med.* **7**, 74 (2024).

[13] Sadek, I., Biswas, J. & Abdulrazak, B. Ballistocardiogram signal processing: a review. *Health Inf. Sci. Syst.* **7**, 10 (2019).

[14] Kim, C.-S. et al. Ballistocardiogram: Mechanism and potential for unobtrusive cardiovascular health monitoring. *Sci. Rep.* **6**, 31297 (2016).

[15] Yu, B., Chen, Y., Cai, D. & Zhang, H. A proof-of-concept study on non-contact bcg-based cardiac monitoring for in-patients with sleep apnea syndrome using piezoelectric ceramics. *IEEE Trans. Instrum. Meas.* (2025).

[16] Paalasmaa, J., Waris, M., Toivonen, H., Leppäkorpi, L. & Partinen, M. *Unobtrusive online monitoring of sleep at home*, 3784–3788 (IEEE, 2012).

[17] Sadek, I., Heng, T. T. S., Seet, E. & Abdulrazak, B. A new approach for detecting sleep apnea using a contactless bed sensor: Comparison study. *J. Med. Internet Res.* **22**, e18297 (2020).

[18] Yi, R., Enayati, M., Keller, J. M., Popescu, M. & Skubic, M. *Non-invasive in-home sleep stage classification using a ballistocardiography bed sensor*, 1–4 (IEEE, 2019).

[19] Migliorini, M. et al. *Automatic sleep staging based on ballistocardiographic signals recorded through bed sensors*, 3273–3276 (IEEE, 2010).

[20] Wiard, R. M. et al. Automatic detection of motion artifacts in the ballistocardiogram measured on a modified bathroom scale. *Med. Biol. Eng. Comput.* **49**, 213–220 (2011).

[21] Vanderperren, K. et al. Removal of bcg artifacts from eeg recordings inside the mr scanner: a comparison of methodological and validation-related aspects. *Neuroimage* **50**, 920–934 (2010).

[22] Mitsukura, Y., Sumali, B., Nagura, M., Fukunaga, K. & Yasui, M. Sleep stage estimation from bed leg ballistocardiogram sensors. *Sensors* **20**, 5688 (2020).

[23] Wang, Z. et al. Assessing the severity of sleep apnea syndrome based on ballistocardiogram. *PLoS One* **12**, e0175351 (2017).

[24] He, K., Zhang, X., Ren, S. & Sun, J. *Deep residual learning for image recognition*, 770–778 (2016).

[25] Su, J. et al. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing* **568**, 127063 (2024).

[26] Korkalainen, H. et al. Deep learning enables sleep staging from photoplethysmogram for patients with suspected sleep apnea. *Sleep* **43**, zsaa098 (2020).

[27] Sun, H. et al. Sleep staging from electrocardiography and respiration with deep learning. *Sleep* **43**, zsz306 (2020).

[28] Sridhar, N. et al. Deep learning for automated sleep staging using instantaneous heart rate. *NPJ Digit. Med.* **3**, 106 (2020).

[29] Radha, M. et al. A deep transfer learning approach for wearable sleep stage classification with photoplethysmography. *NPJ Digit. Med.* **4**, 135 (2021).

[30] Li, S. et al. Sleepnetzero: Zero-burden zero-shot reliable sleep staging with neural networks based on ballistocardiograms. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* **8**, 1–25 (2024).

[31] Liu, M.-H. et al. Efficientnet-based machine learning architecture for sleep apnea identification in clinical single-lead ecg signal data sets. *BioMed. Eng. OnLine* **23**, 57 (2024).

[32] Zhang, Y. et al. Deep learning for obstructive sleep apnea detection and severity assessment: A multimodal signals fusion multiscale transformer model. *Nat. Sci. Sleep* 1–15 (2025).

[33] He, H. et al. What radio waves tell us about sleep! *Sleep* **48**, zsae187 (2025).

[34] Van Pee, B. et al. A multicentric validation study of a novel home sleep apnea test based on peripheral arterial tonometry. *Sleep* **45**, zsac028 (2022).

[35] Raphelson, J. R. et al. Evaluation of a novel device to assess obstructive sleep apnea and body position. *J. Clin. Sleep Med.* **19**, 1643–1649 (2023).

[36] Goetz, L., Seedat, N., Vandersluis, R. & van der Schaar, M. Generalization—a key challenge for responsible ai in patient-facing clinical applications. *NPJ Digit. Med.* **7**, 126 (2024).

[37] Cohen, J. P. et al. Problems in the deployment of machine-learned models in health care. *Cmaj* **193**, E1391–E1394 (2021).

[38] Sadek, I., Heng, T. T. S., Seet, E. & Abdulrazak, B. A new approach for detecting sleep apnea using a contactless bed sensor: Comparison study. *J. Med. Internet Res.* **22**, e18297 (2020).

[39] Chen, X. et al. *Non-contact detection of sleep apnea using multi-timescale cardiopulmonary features*, 439–446 (2024).

[40] Sanchez Gomez, J., Pramono, R. X. A., Imtiaz, S. A., Rodriguez-Villegas, E. & Valido Morales, A. Validation of a wearable medical device for automatic diagnosis of osa against standard psg. *J. Clin. Med.* **13**, 571 (2024).

[41] Davies, C. et al. A single-arm, open-label, multicenter, and comparative study of the anne sleep system vs polysomnography to diagnose obstructive sleep apnea. *J. Clin. Sleep Med.* **18**, 2703–2712 (2022).

[42] Pépin, J.-L. et al. Assessment of mandibular movement monitoring with machine learning analysis for the diagnosis of obstructive sleep apnea. *JAMA Netw. Open* **3**, e1919657–e1919657 (2020).

[43] Jaworski, D. & Park, E. J. Nonlinear heart rate variability analysis for sleep stage classification using integration of ballistocardiogram and apple watch. *Nat. Sci. Sleep* 1075–1090 (2024).

[44] Kwon, H. B. et al. Attention-based lstm for non-contact sleep stage classification using ir-uwb radar. *IEEE J. Biomed. Health Inf.* **25**, 3844–3853 (2021).

[45] Pini, N. et al. An automated heart rate-based algorithm for sleep stage classification: Validation using conventional polysomnography and an innovative wearable electrocardiogram device. *Front. Neurosci.* **16**, 974192 (2022).

[46] Milton, S. et al. Five-year changes in 24-hour sleep-wake activity and dementia risk in oldest old women. *Neurology* **104**, e213403 (2025).

[47] Li, P. et al. Daytime napping and alzheimer's dementia: a potential bidirectional relationship. *Alzheimers Dement.* **19**, 158–168 (2023).

[48] Rosenberg, R. S. & Van Hout, S. The american academy of sleep medicine inter-scorer reliability program: sleep stage scoring. *J. Clin. Sleep Med.* **9**, 81–87 (2013).

[49] Thomson, W. Delay networks having maximally flat frequency characteristics. *Proc. IEE-Part III: Radio Commun. Eng.* **96**, 487–490 (1949).

[50] Brüser, C., Winter, S. & Leonhardt, S. Robust inter-beat interval estimation in cardiac vibration signals. *Physiol. Meas.* **34**, 123 (2013).

[51] Makowski, D. et al. Neurokit2: A python toolbox for neurophysiological signal processing. *Behav. Res. Methods* 1–8 (2021).

[52] Massaro, S. & Pecchia, L. Heart rate variability (hrv) analysis: A methodology for organizational neuroscience. *Organ. Res. Methods* **22**, 354–393 (2019).

[53] Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollar, P. *Focal loss for dense object detection* (2017).

Acknowledgments

The authors gratefully acknowledge the National Sleep Research Resource (NSRR; sleepdata.org), funded by the National Heart, Lung, and Blood Institute (NHLBI), for providing access to de-identified clinical data that supported the development and evaluation of our work. We also thank the Brain Data Science Platform (BDSP) for access to the Human Sleep Project (HSP) dataset, a growing collection of clinical polysomnography recordings from Massachusetts General Hospital. This work is supported by the Ministry of Science and Technology of China STI2030-Major Projects (No. 2021ZD0201900, 2021ZD0201902), and the National Natural Science Foundation of China (62102008).

Author contributions

Conception and design of the study: S.C., X.C., H.Y., X.T., Y.L.; Project supervision and administration: R.Z., D.L., C.L., Y.L., B.W.; Provision of clinical resources: W.H., Y.S., F.L., X.X.; Algorithm design and development: S.C., X.C., C.S., X.J., Y.W., C.Z.; Manuscript drafting: S.C., X.C., S.H., Y.L.; All authors revised the manuscript and approved the final version.

Competing interests

The authors declare no competing interests.