Omni-Sleep: A Sleep Foundation Model via Hierarchical Contrastive Learning of CNS–ANS Dynamics

Zhoujie Hou[1,3*], Song Wang[1,3*], Kexin Lou[1,3*], Mo Wang[1,3], Chen Wei[1,3], and Quanying Liu[1,2,3✉]

[1] Department of Biomedical Engineering, Southern University of Science and Technology, Shenzhen, China
{houzj,12332575,12063004,12250099,weic3}@mail.sustech.edu.cn
liuqy@sustech.edu.cn
[2] Shenzhen Loop Area Institute, Shenzhen, China
[3] Omni-Intelligence, Shenzhen, China

Abstract. Sleep physiology arises from the coordinated dynamics of the central nervous system (CNS) and autonomic nervous system (ANS), as reflected by multimodal polysomnography signals including EEG, EOG, EMG, ECG, and respiration. However, existing sleep foundation models often fuse heterogeneous biosignals in a topology-agnostic manner, overlooking their physiological organization. We introduce Omni-Sleep, a sleep foundation model that uses the CNS/ANS partition as a physiological prior for topology-constrained representation learning. Omni-Sleep learns structured representations through three objectives: intra-system consistency, which captures shared subsystem-level factors within neural and cardio-respiratory signals; inter-system synchronization, which aligns subsystem trajectories to model brain–body dynamics; and latent-space masked temporal modeling, which captures long-horizon sleep dynamics. Pre-trained on over 100,000 hours of multi-center multimodal PSG data, Omni-Sleep is evaluated on sleep staging and multi-disease classification. Across datasets and modality-ablation settings, Omni-Sleep outperforms strong foundation-model baselines, showing improved label efficiency, cross-dataset generalization, and robustness to missing modalities. These results highlight the value of physiological hierarchy for generalizable sleep representation learning. Code is available at `https://github.com/AutoBrain-sleep/OmniSleep`.

Keywords: Sleep · Self-supervised learning · Multimodal.

1 Introduction

Sleep is a dynamic physiological process governed by coordinated interactions between central neural activity and autonomic cardio-respiratory regulation [19,2].

* Equal contribution.
✉ Corresponding author.

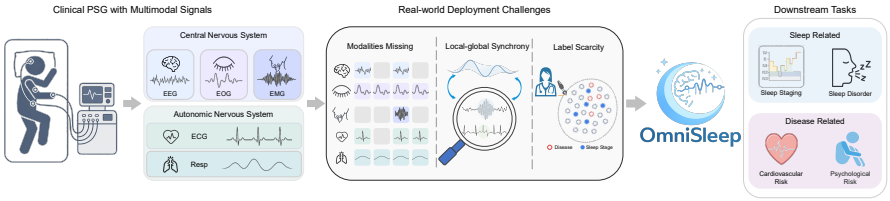

Fig. 1. Omni-Sleep targets real-world PSG deployment.

In clinical practice, this brain–body coupling is captured by polysomnography (PSG), which records heterogeneous signals spanning EEG, EOG, EMG, ECG and respiratory channels. These signals are informative not only for sleep staging, but also for assessing multi-system health risks that manifest through altered sleep physiology [13,1,17,5,8]. Recent self-supervised foundation models for neural data, have shown that large-scale pretraining can learn transferable brain representations and reduce reliance on costly expert labels [10,22,4,21,23,20,24]. This paradigm has also been adopted for sleep analysis through large-scale pretraining on polysomnographic recordings [6,17]. However, existing multimodal methods often align heterogeneous physiological signals in a unified representation space, overlooking their physiological organization [12,17,14,16]. This flat fusion neglects a key property of PSG: CNS-derived signals and ANS-related signals follow distinct physiological manifolds, yet exhibit stage-dependent and time-varying synchronization [25,27]. In real-world deployments, these limitations are amplified under domain shifts and missing modalities. [4].

To bridge this gap, we propose Omni-Sleep, a sleep foundation model that uses the CNS/ANS partition as a physiological prior for topology-constrained sleep representation learning. Pretrained on over 100,000 hours of multi-center PSG data, Omni-Sleep combines hierarchical contrastive learning to capture intra-system consistency and inter-system coupling with long-horizon latent prediction to model macro-scale sleep dynamics. Across multiple independent cohorts, Omni-Sleep demonstrates strong generalization and robustness in sleep staging and sleep-disorder assessment under label scarcity and missing-modality settings. The main contributions of this work are summarized as follows:

- **Structure-Aware Foundation Model:** We propose Omni-Sleep, a sleep foundation model that uses the CNS/ANS split as a physiological prior while learning subsystem-level and cross-system representations.
- **Hierarchical Pretraining framework:** Omni-Sleep combines intra-system contrastive learning, inter-system alignment, and long-horizon latent prediction to capture structured CNS–ANS sleep dynamics.
- **SOTA Performance across Tasks:** Across datasets and modality-ablation settings, Omni-Sleep consistently improves out-of-domain transfer, label efficiency, and robustness to missing modalities, outperforming state-of-the-art baselines in sleep staging and multi-disease classification.

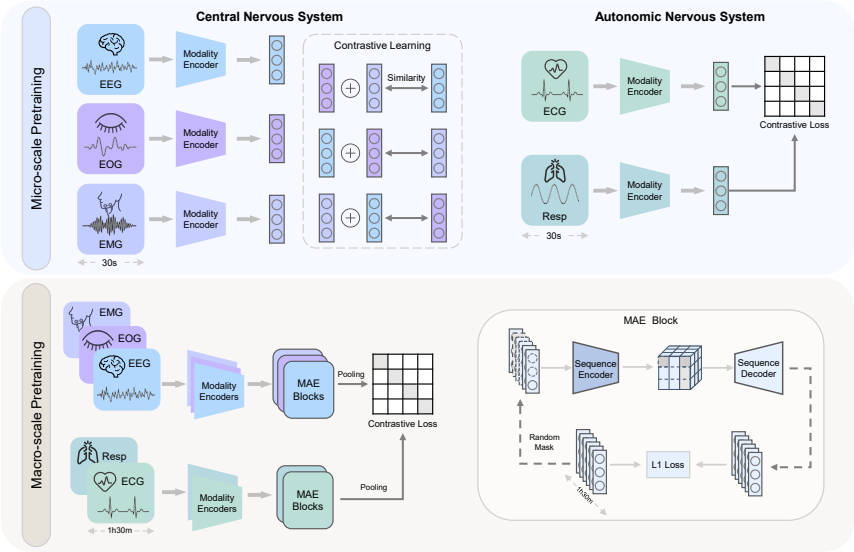

Fig. 2. Omni-Sleep pretraining. **Micro-level** objectives capture intra-system consistency within CNS and ANS signals, while **macro-level** objectives model inter-system synchronization and long-horizon sleep dynamics.

2 Methodology

2.1 Problem Formulation

PSG comprises a set of synchronized physiological modalities $\mathbf{X} = \{X^{(m)}\}_{m\in\mathcal{M}}$, where each $X^{(m)}$ is segmented into a sequence of 30 s epochs. We explicitly encode the physiological information $\mathcal{M}$ by partitioning modalities into CNS subset $\mathcal{M}_C$ (EEG, EOG and EMG) and ANS subset $\mathcal{M}_A$ (ECG and respiration) with $\mathcal{M} = \mathcal{M}_C \cup \mathcal{M}_A$. Our goal is to pretrain a multimodal encoder that preserves both within-group consistency and cross-group synchrony.

2.2 Modality encoders

For each modality m, we transform the raw epoch signal $X^{(m)}$ into a sequence of latent tokens using a lightweight 1D convolutional patch encoder:

$$\mathbf{H}^{(m)} = f^{(m)}(\mathbf{X}^{(m)}), \qquad \mathbf{H}^{(m)} \in \mathbb{R}^{T\times D}, \tag{1}$$

where T is the token length and D is the latent dimension.

To model long-range temporal dependencies, tokens are passed through a modality-specific sequence encoder $g^{(m)}_{\text{seq}}$ (RoFormer with rotary positional encoding), providing contextualized tokens $\mathbf{S}^{(m)} = g^{(m)}_{\text{seq}}(\mathbf{H}^{(m)})$, $\mathbf{S}^{(m)} \in \mathbb{R}^{T\times D}$ and

an epoch-level embedding obtained by pooling:

$$\mathbf{s}^{(m)} = \mathrm{Pool}\Big(S^{(m)}\Big) \in \mathbb{R}^D. \tag{2}$$

We denote the projected representation for contrastive learning as $\mathbf{z}^{(m)} = \psi(\mathbf{s}^{(m)})$, where $\psi(\cdot)$ is a small MLP projection head.

2.3 Hierarchical Contrastive Learning

Omni-Sleep optimizes a hierarchical contrastive objective that (i) enforces **intra-system consistency** within CNS and within ANS (micro-scale), and (ii) preserves **inter-system coupling** between pooled CNS and pooled ANS summaries (macro-scale).

InfoNCE objective. Given an anchor representation $\mathbf{z}_i$ and its positive $\mathbf{z}_i^+$, we use standard InfoNCE:

$$\ell_{\mathrm{NCE}}(\mathbf{z}_i, \mathbf{z}_i^+) = -\log \frac{\exp(\mathrm{sim}(\mathbf{z}_i, \mathbf{z}_i^+)/\tau)}{\sum_{j=1}^{B} \exp(\mathrm{sim}(\mathbf{z}_i, \mathbf{z}_j)/\tau)}, \tag{3}$$

where $\mathrm{sim}(\cdot,\cdot)$ is cosine similarity, τ is a temperature, B is the batch size, and negatives are other samples in the mini-batch.

Micro-scale (intra-system) contrastive learning. For each modality m within a subsystem $\mathcal{S} \in \{\mathcal{M}_C, \mathcal{M}_A\}$, we construct the positive as the mean of the other modalities in the same subsystem and same epoch:

$$\mathbf{z}^{(m)+} = \frac{1}{|\mathcal{S}| - 1} \sum_{m' \in \mathcal{S} \setminus \{m\}} \mathbf{z}^{(m')}. \tag{4}$$

The intra-system loss is

$$\mathcal{L}_{\mathrm{intra}} = \sum_{\mathcal{S} \in \{\mathcal{M}_C, \mathcal{M}_A\}} \sum_{m \in \mathcal{S}} \ell_{\mathrm{NCE}}\Big(\mathbf{z}^{(m)}, \mathbf{z}^{(m)+}\Big). \tag{5}$$

This encourages each modality to agree with the shared subsystem representation, improving robustness when some channels are absent.

Macro-scale (inter-system) alignment. To explicitly model brain–body coupling, we form pooled subsystem summaries

$$\mathbf{z}_C = \frac{1}{|\mathcal{M}_C|} \sum_{m \in \mathcal{M}_C} \mathbf{z}^{(m)}, \qquad \mathbf{z}_A = \frac{1}{|\mathcal{M}_A|} \sum_{m \in \mathcal{M}_A} \mathbf{z}^{(m)}. \tag{6}$$

We then align CNS and ANS summaries using a symmetric InfoNCE:

$$\mathcal{L}_{\mathrm{inter}} = \ell_{\mathrm{NCE}}(\mathbf{z}_C, \mathbf{z}_A) + \ell_{\mathrm{NCE}}(\mathbf{z}_A, \mathbf{z}_C). \tag{7}$$

Contrastive objective. The topology-aware contrastive loss is

$$\mathcal{L}_c = \mathcal{L}_{\text{intra}} + \lambda \, \mathcal{L}_{\text{inter}}, \tag{8}$$

where λ controls the strength of CNS–ANS coupling.

2.4 Long-horizon latent masked modeling (macro-scale)

Local morphology within a single epoch is insufficient to capture long-range sleep dynamics. We therefore introduce a long-horizon masked prediction objective over windows comprising multiple consecutive epochs (e.g., an extended context such as 1.5 h). For each modality m, we collect the epoch embeddings in a window:

$$\mathbf{E}^{(m)} = [\mathbf{h}_1^{(m)}, \ldots, \mathbf{h}_L^{(m)}] \in \mathbb{R}^{L \times D}, \tag{9}$$

where L is the number of epochs in the long window. We randomly mask a subset of epoch positions Ω and feed the unmasked sequence to a temporal RoFormer to predict embeddings at masked positions:

$$\hat{\mathbf{h}}_t^{(m)} = g_{\text{pred}}^{(m)}\left(\mathbf{E}_{\setminus \Omega}^{(m)}\right), \quad t \in \Omega. \tag{10}$$

We use an ℓ_1 reconstruction loss to encourage robust, long-horizon representations:

$$\mathcal{L}_p = \sum_{m \in \mathcal{M}} \sum_{t \in \Omega} \left\| \mathbf{h}_t^{(m)} - \hat{\mathbf{h}}_t^{(m)} \right\|_1. \tag{11}$$

This objective complements $\mathcal{L}_c$ by enforcing consistency over extended temporal context, improving transfer under domain shift and incomplete observations.

2.5 Overall training objective and schedule

The final pretraining loss is

$$\mathcal{L} = \alpha \, \mathcal{L}_c + \beta \, \mathcal{L}_p, \tag{12}$$

with weights α and β. We use a two-stage schedule: (i) warm-up pretraining with topology-aware contrastive learning ($\mathcal{L}_c$), followed by (ii) joint optimization of $\mathcal{L}_c$ and long-horizon masked prediction ($\mathcal{L}_p$).

2.6 Implementation details

Omni-Sleep contains about **56 M parameters** in total. Pretraining is conducted on a single NVIDIA A800 GPU for 40 epochs, with ∼40 minutes per epoch. We optimize using AdamW with a cosine learning-rate schedule.

3 Experiments

To evaluate the efficacy of Omni-Sleep in learning unified and robust representations, we designed an experimental protocol spanning multiple scales of data availability, domain shifts, and clinical pathologies.

3.1 Experimental Setup

Datasets. Our pre-training corpus comprises over 100,000 hours of multi-center PSG recordings from SHHS [15], WSC [26], and MESA [3]. A random subset of 1,159 SHHS subjects was strictly held out for downstream disease classification. For independent out-of-domain (OOD) evaluation, we utilized ISRUC-Sleep Subgroup I (100 subjects)[11] and CinC (994 subjects) [7], ensuring absolute separation from the pre-training phase.All PSG recordings underwent subject-specific normalization. To efficiently standardize multimodal inputs, high-frequency channels (EEG, EOG, EMG, ECG) were uniformly resampled to 100 Hz, while lower-frequency respiratory and airflow channels were resampled to 10 Hz.

Implementation Details. Omni-Sleep is implemented in PyTorch and optimized with AdamW using a cosine learning rate schedule. We compare Omni-Sleep (56M parameters) with SleepFM (4.4M) and SleepGPT (134M) as baselines. Representations are evaluated using both linear probing with a frozen backbone and full fine-tuning in few-shot settings. Performance is measured by Macro F1-score, Cohen's Kappa (κ), and Accuracy for sleep staging, and by AUROC for disease classification.

Table 1. Linear-probing performance on sleep staging across PSG channel sets. We evaluate frozen representations using different channel configurations. **Omni-Sleep (ours)** denotes the final proposed strategy. CNS and ANS refer to central and autonomic nervous system signal sets, respectively. Best results are highlighted in **bold**.

Modality	Model	Overall (↑)			Class-wise F1 (↑)				
		Acc.	κ	MF1	W	N1	N2	N3	REM
EEG	SleepFM [17]	65.8	0.54	64.3	76.1	37.8	74.2	73.4	59.9
	SleepGPT [9]	68.6	0.59	67.2	79.7	45.1	72.3	75.4	63.6
	Omni-Sleep (ours)	**74.0**	**0.66**	**73.5**	**82.7**	**56.4**	**75.7**	**77.1**	**75.6**
CNS	SleepFM [17]	70.5	0.61	69.1	78.8	42.6	78.9	74.1	71.1
	SleepGPT [9]	70.8	0.62	69.9	80.3	48.8	73.9	76.3	70.3
	Omni-Sleep (ours)	**76.7**	**0.70**	**76.2**	**84.2**	**59.6**	**78.4**	**78.5**	**80.3**
ANS	SleepFM [17]	58.4	0.42	57.1	62.5	30.2	69.8	71.5	51.5
	Omni-Sleep (ours)	**66.3**	**0.56**	**65.3**	**71.5**	**44.8**	**69.7**	**74.7**	**66.1**
Full	SleepFM [17]	71.8	0.63	70.0	79.0	43.9	79.3	74.3	71.7
	SleepGPT [9]	70.8	0.62	69.9	80.3	48.8	73.9	76.3	70.3
	Omni-Sleep (ours)	**77.8**	**0.71**	**77.3**	**85.1**	**60.6**	**79.5**	**78.9**	**82.4**

3.2 Out-of-Domain Generalization

To assess intrinsic feature quality and transferability, we evaluated Omni-Sleep on the external CinC 2018 dataset under strict linear probing. By explicitly modeling the distinct manifolds of the CNS and ANS during pre-training, Omni-Sleep

extracts invariant features that generalize to unseen clinical environments, outperforming existing flat-fusion foundation models without task-specific weight updates. As detailed in Table 1, Omni-Sleep consistently establishes new state-of-the-art benchmarks across all evaluated channel combinations. In the Full modality setting, our model achieves an Accuracy of 77.8% and a Macro F1 of 77.3%, yielding substantial absolute improvements over existing models. Crucially, Omni-Sleep demonstrates remarkable resilience to sensor dropouts; even in restricted scenarios such as ANS-only or EEG-only, it substantially exceeds the performance of both SleepFM and SleepGPT. This validates that our hierarchical alignment strategy effectively extracts generalized and robust representations independent of sensor availability.

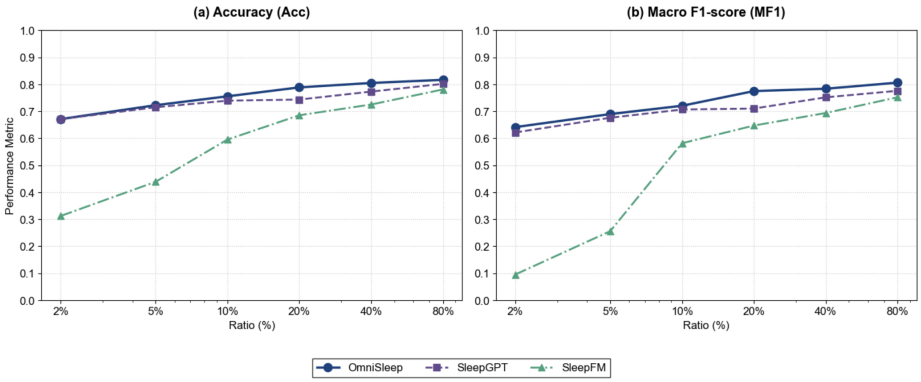

Fig. 3. Few-shot sleep staging on ISRUC-I.

3.3 Label Efficiency in Clinical Scenarios

To evaluate label efficiency under realistic low-annotation settings, we fine-tuned the model using 2%—80% labeled data. As shown in Fig. 3, Omni-Sleep consistently outperformed SleepGPT and SleepFM across all ratios, remaining competitive even at 2% supervision and progressively approaching the fully supervised regime as label availability increased. These findings indicate that the proposed temporal predictive pre-training effectively captures clinically meaningful transition dynamics, thereby reducing downstream annotation requirements.

3.4 Multi-Label Disease Classification

We evaluated Omni-Sleep for multi-disease classification on SHHS1 across sleep-related, psychiatric, respiratory, and cardiovascular conditions. Binary labels were defined using established clinical criteria, including AHI > 5 for OSA, ESS ≥ 12 for hypersomnia, MCS ≤ 42 for clinical depression, and "So Blue"

score ≥ 3 for depressive affect. Respiratory outcomes, including COPD and emphysema, and cardiovascular outcomes, including heart failure and stroke, were derived from physician-confirmed diagnoses. As shown in Table 2, linear probing on frozen Omni-Sleep representations consistently outperforms demographic baselines and prior foundation models. The ablation variant Omni-Sleep* further suggests that long-horizon temporal modeling contributes to classification of respiratory and cardiovascular conditions, while modality-aggregation results demonstrate robustness to partial observability and sensor dropouts.

Table 2. Performance comparisons on multi-label disease classification. We report AUROC across disease classification tasks via linear probing. Values are reported as mean ± standard deviation over five runs, with the best results highlighted in **bold**. Omni-Sleep* denotes an ablation variant without the RoFormer module.

Method	Sleep Disorders		Respiratory Disorders	
	OSA	Hypersomnia	COPD	Emphysema
Demographic [18]	0.774±0.017	0.423±0.328	0.562±0.175	0.632±0.201
SleepGPT [9]	0.597±0.029	0.460±0.123	0.626±0.119	0.649±0.071
Omni-Sleep*	0.789±0.044	0.471±0.053	0.665±0.191	0.692±0.060
Omni-Sleep (ours)	**0.841±0.012**	**0.551±0.109**	**0.750±0.158**	**0.750±0.043**

Method	Psychological		Cardiovascular	
	Clin. Depression	Depressive Affect	Heart Fail	Stroke
Demographic [18]	0.513±0.037	0.545±0.124	0.625±0.024	0.707±0.144
SleepGPT [9]	0.516±0.078	0.583±0.091	0.409±0.147	0.507±0.043
Omni-Sleep*	**0.545±0.057**	0.565±0.053	0.796±0.063	0.735±0.078
Omni-Sleep (ours)	0.519±0.051	**0.623±0.041**	**0.825±0.066**	**0.739±0.078**

4 Conclusion

We propose Omni-Sleep, a novel multimodal foundation model that uses the CNS/ANS partition as a physiological prior for topology-constrained sleep representation learning. Pre-trained on over 100,000 hours of multimodal PSG data, Omni-Sleep combines intra-system consistency, inter-system synchronization, and long-horizon latent prediction to capture subsystem-level factors, CNS–ANS coupling, and macro-scale sleep dynamics. Across datasets and modality-ablation settings, Omni-Sleep improves generalization, few-shot label efficiency, and missing-modality robustness for sleep staging and sleep-related disorder classification, highlighting the value of physiological hierarchy for sleep representation learning.

Acknowledgments. This work was supported in part by the National Natural Science Foundation of China, under Grant Nos. 3254100307 and 62472206; the Brain Science

and Brain-like Intelligence Technology National Science and Technology Major Project, under Grant No. 2021ZD0200500; the National Key R&D Program of China, under Grant No. 2025YFC3410000; and the Shenzhen Science and Technology Innovation Committee, under Grant Nos. RCYX20231211090405003 and JCYJ20220818100213029. Additional support was provided by the Guangdong Basic and Applied Basic Research Foundation, under Grant No. 2026B1515020099; the Guangdong S&T Program, under Grant No. 2026B0101110003; the Shanghai Municipal Special Program for Basic Research on General AI Foundation Models, under Grant No. 2025SHZDZX026D05; the Shenzhen Loop Area Institute, under Grant No. FPF10120250012; the open research fund of the Guangdong Provincial Key Laboratory of Mathematical and Neural Dynamical Systems; the Center for Computational Science and Engineering at Southern University of Science and Technology; and the Shenzhen Key Laboratory of Smart Healthcare Engineering.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdalbari, H., Durrani, M., Pancholi, S., Patel, N., Nasuto, S.J., Nicolaou, N.: Brain and brain-heart granger causality during wakefulness and sleep. Frontiers in Neuroscience **16**, 927111 (2022)
2. Bashan, A., Bartsch, R.P., Kantelhardt, J.W., Havlin, S., Ivanov, P.C.: Network physiology reveals relations between network topology and physiological function. Nature communications **3**(1), 702 (2012)
3. Chen, X., Wang, R., Zee, P., Lutsey, P.L., Javaheri, S., Alcántara, C., Jackson, C.L., Williams, M.A., Redline, S.: Racial/ethnic differences in sleep disturbances: the multi-ethnic study of atherosclerosis (mesa). Sleep **38**(6), 877–888 (2015)
4. Chien, H.Y.S., Goh, H., Sandino, C.M., Cheng, J.Y.: Maeeg: Masked auto-encoder for eeg representation learning. arXiv preprint arXiv:2211.02625 (2022)
5. Fox, B., Jiang, J., Wickramaratne, S.D., Kovatch, P., Suarez-Farinas, M., Shah, N.A., Parekh, A., Nadkarni, G.N.: A foundational transformer leveraging full night, multichannel sleep study data accurately classifies sleep stages. Sleep (2025), https://api.semanticscholar.org/CorpusID:276988590
6. Gerardy, B., Kuna, S.T., Pack, A., Kushida, C.A., Walsh, J.K., Staley, B., Pien, G.W., Younes, M.: An approach for determining the reliability of manual and digital scoring of sleep stages. Sleep **46**(11), zsad248 (2023)
7. Ghassemi, M.M., Moody, B.E., Lehman, L.w.H., Song, C., Li, Q., Sun, H., Westover, M.B., Clifford, G.D.: You snooze, you win: the physionet/computing in cardiology challenge 2018. In: 2018 Computing in Cardiology Conference (CinC). vol. 45, pp. 1–4. IEEE (2018). https://doi.org/10.22489/CinC.2018.049
8. Goldammer, M., Zaunseder, S., Brandt, M.D., Malberg, H., Gräßer, F.: Investigation of automated sleep staging from cardiorespiratory signals regarding clinical applicability and robustness. Biomedical Signal Processing and Control **71**, 103047 (2022)
9. Huang, W., Wang, Y., Cheng, H., Xu, W., Li, T., Wu, X., Xu, H., Liao, P., Cui, Z., Zou, Q., Gao, J.: A unified time-frequency foundation model for sleep decoding. Nature Communications **17** (2026), https://api.semanticscholar.org/CorpusID:284715730

10. Jiang, W.B., Zhao, L., Lu, B.L.: Large brain model for learning generic representations with tremendous eeg data in bci. In: International Conference on Learning Representations. vol. 2024, pp. 16405–16426 (2024)
11. Khalighi, S., Sousa, T., Santos, J.M., Nunes, U.: Isruc-sleep: A comprehensive public dataset for sleep researchers. Computer methods and programs in biomedicine **124**, 180–192 (2016)
12. Liu, P., Qian, W., Zhang, H., Zhu, Y., Hong, Q., Li, Q., Yao, Y.: Automatic sleep stage classification using deep learning: signals, data representation, and neural networks. Artificial Intelligence Review **57**(11), 301 (2024)
13. Lyu, J., Yeh, C.H., Shi, W., Penzel, T.: Dynamic brain-heart interaction in sleep characterized by variational phase-amplitude coupling framework. Communications Biology **8**(1), 1235 (2025)
14. Phan, H., Chén, O.Y., Tran, M.C., Koch, P., Mertins, A., De Vos, M.: Xsleepnet: Multi-view sequential model for automatic sleep staging. IEEE Transactions on Pattern Analysis and Machine Intelligence **44**(9), 5903–5915 (2021)
15. Quan, S.F., Howard, B.V., Iber, C., Kiley, J.P., Nieto, F.J., O'Connor, G.T., Rapoport, D.M., Redline, S., Robbins, J., Samet, J.M., et al.: The sleep heart health study: design, rationale, and methods. Sleep **20**(12), 1077–1085 (1997)
16. Seo, H., Back, S., Lee, S., Park, D., Kim, T., Lee, K.: Intra-and inter-epoch temporal context network (iitnet) using sub-epoch features for automatic sleep scoring on raw single-channel eeg. Biomedical signal processing and control **61**, 102037 (2020)
17. Thapa, R., He, B., Kjaer, M.R., Moore IV, H., Ganjoo, G., Mignot, E., Zou, J.Y.: Sleepfm: Multi-modal representation learning for sleep across ecg, eeg and respiratory signals. In: AAAI 2024 Spring Symposium on Clinical Foundation Models (2024)
18. Thapa, R., Kjaer, M.R., He, B., Covert, I., HyattMoore, I., Hanif, U., Ganjoo, G., Westover, M.B., Jennum, P., Brink-Kjaer, A., Mignot, E., Zou, J., Ruud, M.R.: A multimodal sleep foundation model for disease prediction. Nature Medicine **32**, 752 – 762 (2026), `https://api.semanticscholar.org/CorpusID:284529751`
19. Thomas, R.J., Mietus, J.E., Peng, C.K., Goldberger, A.L.: An electrocardiogram-based technique to assess cardiopulmonary coupling during sleep. Sleep **28**(9), 1151–1161 (2005)
20. Wang, M., Xia, J., Ye, W., Liu, E., Peng, K., Feng, J., Liu, Q., Wen, H.: Slimbrain: A data-and training-efficient foundation model for fmri data analysis. arXiv preprint arXiv:2512.21881 (2025)
21. Wang, M., Ye, W., Xia, J., Xu, M., Wen, H., Liu, Q.: Flexibrain: Resolution-agnostic voxel-level encoding for native fmri. arXiv preprint arXiv:2606.11500 (2026)
22. Wang, M., Ye, W., Xia, J., Zhang, J., Pan, X., Xu, M., Deng, H., Wen, H., Liu, Q.: Omni-fmri: A universal atlas-free fmri foundation model. arXiv preprint arXiv:2601.23090 (2026)
23. Xia, J., Ye, W., Pan, X., Shen, X., Wang, M., Liu, Q.: Brain-dit: A universal multi-state fmri foundation model with metadata-conditioned pretraining. arXiv preprint arXiv:2604.12683 (2026)
24. Xia, J., Ye, W., Zhang, J., Pan, X., Wang, M., Liu, Q.: Brainworld: A structural-prior-conditioned generative model for whole-brain 4d fmri dynamics. arXiv preprint arXiv:2606.17742 (2026)
25. Yetton, B.D., McDevitt, E.A., Cellini, N., Shelton, C., Mednick, S.C.: Quantifying sleep architecture dynamics and individual differences using big data and bayesian networks. PloS one **13**(4), e0194604 (2018)

26. Young, T.: Rationale, design, and findings from the wisconsin sleep cohort study: toward understanding the total societal burden of sleep-disordered breathing. Sleep medicine clinics **4**(1), 37–46 (2009)
27. de Zambotti, M., Trinder, J., Silvani, A., Colrain, I.M., Baker, F.C.: Dynamic coupling between the central and autonomic nervous systems during sleep: a review. Neuroscience & Biobehavioral Reviews **90**, 84–103 (2018)