

# Deep Learning-Based Prediction of Neural Ionic Homeostasis Changes During Sleep Pressure Accumulation

Huijia Gao

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.  
hellohuijia@colostate.edu

Jeffrey Taylor

Department of Computer Science, Binghamton University, Binghamton, NY, USA.  
jeffrey.taylor@binghamton.edu

Xiaohai Chen

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV,  
USA.  
xchen@unr.edu

## Abstract

Sleep pressure accumulation is a fundamental physiological process that drives the transition from wakefulness to sleep, yet its underlying mechanisms at the ionic and neural circuit levels remain incompletely understood. Recent advances in genetically encoded biosensors have enabled real-time monitoring of intracellular and extracellular ion concentrations, revealing dynamic shifts in neural ionic homeostasis during prolonged wakefulness. This paper presents a comprehensive system-level framework for predicting these ionic homeostasis changes using deep learning architectures. We propose an integrated infrastructure that combines high-throughput electrophysiological recordings, calcium and pH imaging data, and molecular biosensor inputs to train predictive models capable of forecasting ion concentration trajectories as sleep pressure builds. The system architecture involves distributed data acquisition from multiple recording modalities, centralized preprocessing pipelines, and scalable model deployment across cloud and edge computing environments. Key design trade-offs are examined regarding model complexity, inference latency, and energy efficiency, particularly for potential integration into wearable sleep-monitoring devices. Robustness and fairness considerations are addressed through stratified data augmentation and adversarial debiasing techniques to mitigate demographic and experimental biases inherent in neurophysiological datasets. Policy and governance implications are discussed with respect to data privacy, algorithmic transparency, and regulatory frameworks for AI-driven neuroscience tools. By bridging molecular neurobiology and deep learning, this work lays a foundation for real-time prediction of sleep pressure dynamics, with broad applications in sleep medicine, cognitive performance optimization, and neural disorder diagnostics.

## Keywords

deep learning, sleep pressure, ionic homeostasis, neural prediction, biosensor data, system architecture, fairness, neurotechnology policy.

## 1. Introduction

Sleep is an essential neurobiological state governed by complex homeostatic and circadian processes. The accumulation of sleep pressure during wakefulness is widely attributed to the gradual buildup of extracellular adenosine, but growing evidence implicates broader ionic shifts, including changes in potassium, sodium, calcium, and proton concentrations, as critical mediators of sleep drive [1, 2]. These ionic fluctuations modulate neuronal excitability, synaptic strength, and network oscillations, ultimately triggering sleep onset when homeostatic thresholds are exceeded. Despite the importance of these dynamics, the ability to predict ionic homeostasis changes in real time has remained elusive due to the high dimensionality of ion regulation pathways, the nonlinear interactions among multiple ion species, and the difficulty of continuous *in vivo* monitoring.

Deep learning offers a powerful approach to model such complex, high-dimensional, and temporally correlated processes. Recurrent neural networks, long short-term memory networks, and transformers have demonstrated remarkable success in time-series forecasting across domains including climate science, finance, and physiology [3, 4]. In neuroscience, deep learning has been applied to decode neural activity patterns from electrophysiological recordings, predict seizure onset, and classify sleep stages from electroencephalography signals [5, 6]. However, the application of deep learning to predict molecular-level ionic homeostasis changes in the brain remains largely unexplored. This paper addresses that gap by proposing a system-level framework for building and deploying deep learning models that can forecast neural ionic concentration changes during sleep pressure accumulation.

The paper is organized as follows. Section 2 reviews the biological foundations of sleep pressure and ionic homeostasis, along with existing computational models and limitations. Section 3 describes the proposed system architecture, including data infrastructure, acquisition modalities, and preprocessing pipelines. Section 4 details the deep learning model design, training strategies, and deployment considerations. Section 5 examines robustness and fairness, highlighting techniques to improve generalization and reduce biases. Section 6 discusses policy and governance implications, focusing on data privacy, algorithmic accountability, and regulatory pathways. Section 7 concludes with a synthesis of findings and future research directions.

## **2. Background and Related Work**

Sleep pressure, also known as sleep drive, increases monotonically during wakefulness and dissipates during sleep. The two-process model posits an interaction between a circadian process and a homeostatic process, the latter reflecting the accumulation of sleep-promoting substances [7]. Adenosine has been the most studied homeostatic factor, but recent work has expanded the repertoire to include ions such as potassium, which accumulates extracellularly during wakefulness and is cleared during sleep [8]. Calcium dynamics, both intracellular and extracellular, also exhibit sleep-wake dependent changes, influencing synaptic plasticity and neuronal firing patterns [9]. Furthermore, intracellular pH fluctuations have been identified as a potential sleep driver, with prolonged wakefulness leading to acidification that is reversed by sleep [10]. The genetically encoded ionic-stress sensor reported in a recent preprint enables direct measurement of proton concentration changes *in vivo*, providing a powerful tool to link ionic homeostasis to sleep pressure [11, 12].

Computational modeling of sleep homeostasis has traditionally relied on differential equations representing concentration changes of key molecules [13]. These models are limited by their reliance on a small number of variables and linear or weakly nonlinear coupling assumptions. In contrast, deep learning approaches can capture high-order interactions and

nonlinear dynamics from multivariate time-series data without explicit mechanistic assumptions. Several studies have used recurrent neural networks to predict sleep stages from electroencephalography and actigraphy data [14, 15]. However, these models operate at the macroscopic level and do not incorporate molecular ionic data. A few recent works have combined calcium imaging with convolutional neural networks to infer neuronal activity states, but they focus on spatial patterns rather than ionic concentration trajectories [16]. The gap between molecular-level ion sensing and predictive deep learning remains substantial, motivating the present system-level design.

The potential benefits of a predictive deep learning framework for ionic homeostasis are significant. Real-time prediction could enable early detection of excessive sleep pressure, inform personalized sleep scheduling, and support closed-loop neuromodulation therapies. Nevertheless, deploying such models in research or clinical settings requires careful consideration of data infrastructure, model architecture, and ethical safeguards. The following sections develop these aspects in depth.

### **3. System Architecture and Data Infrastructure**

The successful prediction of neural ionic homeostasis changes depends critically on a robust data infrastructure that can acquire, store, and preprocess multimodal time-series data from diverse sources. The proposed architecture is designed as a distributed system with three primary layers: the acquisition layer, the preprocessing and storage layer, and the model training and inference layer.

The acquisition layer comprises multiple recording modalities. Electrophysiological recordings using multi-electrode arrays provide high-temporal-resolution measurements of local field potentials and spike activity, which correlate with extracellular ion concentrations [17]. Calcium imaging using genetically encoded indicators such as GCaMP offers indirect readouts of neuronal activity and can be combined with ratiometric pH sensors to estimate proton concentrations [12]. Direct ion-selective microelectrodes can measure potassium and sodium levels with high specificity but limited spatial coverage. A key challenge is synchronizing these heterogeneous data streams, which operate at different sampling rates, spatial resolutions, and noise characteristics. The system adopts a time-stamped unified data format with interpolation algorithms to align measurements at a common temporal grid, typically one hertz or higher depending on the fastest recording modality.

The preprocessing and storage layer handles noise reduction, artifact removal, and feature extraction. For electrophysiology data, bandpass filtering and spike sorting are applied using established pipelines [18]. Calcium imaging data undergo motion correction, region-of-interest segmentation, and deconvolution to estimate spike rates. Ion-selective electrode signals are low-pass filtered to remove high-frequency noise while preserving slow ionic drift. All preprocessed data are stored in a distributed database with metadata describing experimental conditions, animal strain, age, sex, and sleep pressure state. This metadata is essential for later stratification and debiasing. Data storage follows a hierarchical schema that separates raw, intermediate, and processed datasets to facilitate reproducibility and incremental model retraining.

The model training and inference layer is built on a hybrid cloud-edge architecture. Training of deep learning models requires substantial computational resources, including graphics processing units or tensor processing units, which are most cost-effectively hosted in cloud data centers. However, real-time inference for wearable or implantable devices demands on-

device computation with low latency and minimal energy consumption. Therefore, the system supports both cloud-based training and edge-based inference, with model compression techniques such as quantization and pruning to reduce the memory footprint and inference time [19]. Data synchronization between edge devices and the cloud is performed during charging or low-activity periods to minimize network bandwidth usage. The architecture also incorporates a feedback loop where edge-inference predictions are periodically uploaded to retrain and refine cloud models, enabling continuous improvement while respecting privacy constraints.

System-level trade-offs are evident in this design. Higher model complexity improves prediction accuracy but increases inference latency and energy consumption. For edge deployment, simpler recurrent architectures or lightweight transformer variants may be preferable, whereas cloud training can exploit deep transformer networks with attention mechanisms that capture long-range temporal dependencies. Another trade-off exists between data granularity and storage cost. The system must balance the retention of high-resolution raw data for future reanalysis against the practical limitations of storage capacity and retrieval speed. In experimental settings, it is common to retain full raw data for a limited period and then compress or downsample after validation.

Case illustrations from recent sleep research demonstrate the feasibility of such infrastructure. For example, multichannel recordings of potassium and calcium dynamics in the mouse brain during prolonged wakefulness have been successfully collected and analyzed [20]. These datasets, while still limited in size, provide proof-of-concept for training initial deep learning models. Scaling to larger cohorts and longer recording durations will require automated data collection and cloud-based processing, as outlined here.

#### **4. Deep Learning Model Design and Deployment**

Predicting ion concentration changes from multimodal time-series data requires a deep learning architecture that can fuse heterogeneous inputs and model temporal dependencies over multiple timescales. The core of the proposed model is a hybrid encoder-decoder framework. The encoder consists of parallel branches for each input modality: a convolutional neural network branch for calcium imaging sequences, a recurrent long short-term memory branch for electrophysiological traces, and a separate branch for direct ion-selective electrode readings. These branches are fused via a cross-modal attention mechanism that learns to weight features from each modality according to their relevance for predicting specific ion species.

The decoder is a temporal transformer that takes the fused embedding and outputs predictions for multiple future time steps, typically up to 30 minutes ahead at one-second resolution. The transformer's self-attention layers capture long-range dependencies that are critical for modeling the slow accumulation of sleep pressure over hours. To handle the nonstationary nature of ionic data, which includes both circadian trends and transient fluctuations, the model incorporates a learnable trend decomposition module that separates baseline drift from fast dynamics. This decomposition is analogous to the additive decomposition used in classical time-series analysis but is learned end-to-end.

Training the model requires a carefully curated dataset with balanced representation of different sleep pressure levels. Because sleep pressure accumulates gradually, most data points correspond to intermediate levels, with fewer examples at extreme low or high sleep pressure. To avoid predictive bias, the training set is augmented using synthetic data

generated by a variational autoencoder trained on real recordings. This augmentation expands the tails of the distribution, improving the model's ability to predict rare events such as sudden ionic shifts that precede sleep onset. Additionally, a curriculum learning strategy is employed, where the model is first trained on short prediction horizons and gradually increased to longer horizons, stabilizing the learning dynamics.

Deployment of the trained model involves several engineering considerations. In a cloud-based deployment, the model receives preprocessed data streams from multiple experimental subjects simultaneously. The inference pipeline must handle variable-length inputs and produce predictions with bounded latency, typically less than 100 milliseconds per batch to enable real-time feedback. In edge deployment on a wearable device, the model is quantized to eight-bit integer precision and pruned to remove redundant connections, reducing the model size from hundreds of megabytes to a few megabytes. This compression incurs a small accuracy penalty, typically one to two percent in mean absolute error, which is acceptable for many applications.

Sustainability is an important aspect of model deployment. Training deep learning models consumes significant energy, particularly for transformer architectures with millions of parameters. To mitigate environmental impact, the system uses energy-aware scheduling that runs training jobs during periods of low grid carbon intensity, and leverages hardware accelerators with high energy efficiency. For edge devices, energy consumption is minimized by using low-power microcontrollers and limiting inference to a duty cycle that matches the dynamics of ionic changes. For example, predictions can be made every ten seconds rather than every second, reducing energy usage by an order of magnitude.

Robustness to sensor failure and data dropout is built into the inference pipeline. The model is trained with random modality masking, meaning that during training, inputs from certain channels are set to zero with some probability. This forces the encoder to rely on available modalities and prevents overfitting to any single sensor. During deployment, if a sensor malfunctions, the model can still produce reasonable predictions using the remaining channels, gracefully degrading rather than failing completely.

## **5. Robustness and Fairness Considerations**

Deep learning models for physiological prediction are susceptible to biases originating from imbalanced training datasets, experimental confounds, and differences in recording conditions across laboratories. For sleep pressure prediction, biases can arise from the overrepresentation of certain animal strains, age groups, or sleep deprivation protocols. If a model is trained predominantly on young male mice subjected to forced sleep deprivation, it may perform poorly on aged female mice or those undergoing natural sleep restriction. Such disparities raise concerns about the generalizability and fairness of the predictive system, especially if it is eventually translated to human applications.

To address robustness, the proposed framework incorporates stratified data sampling during training. The metadata associated with each recording session (species, sex, age, genetics, and protocol type) is used to define strata, and the training set is constructed to ensure that each stratum is proportionally represented. In cases where some strata have very few samples, synthetic data augmentation is employed, as described earlier. Additionally, adversarial debiasing is implemented by training a discriminator network that attempts to predict the stratum label from the model's learned representations. The main model is simultaneously optimized to minimize the discriminator's accuracy, effectively removing stratum-specific

information from the feature space. This technique has been shown to reduce performance disparities across demographic groups in medical imaging tasks [21] and is adapted here for time-series data.

Another source of bias is the difference in recording equipment and calibration across labs. Even with standardized protocols, minor variations in electrode impedance, filter settings, and baseline drift can introduce systematic offsets that cause a model trained on one dataset to fail on another. The system addresses this through domain adaptation techniques. A small amount of target domain data is used to fine-tune the model's batch normalization layers and a shallow adaptation network, aligning the feature distributions without full retraining [22]. This approach balances the need for generalizability with the practical constraint that collecting large target domain datasets may be expensive.

Fairness also extends to the interpretation of model predictions. When a deep learning model predicts an impending ionic homeostasis shift that could indicate excessive sleep pressure, it is important that the confidence of the prediction is communicated alongside the output. The model can be augmented with a Bayesian layer that produces uncertainty estimates for each prediction. These estimates allow users to distinguish between high-certainty and low-certainty forecasts, which is particularly important when making decisions about intervention, such as administering a wakefulness-promoting drug or initiating a sleep period.

Finally, the system must guard against adversarial attacks. Malicious inputs designed to manipulate the prediction could have serious consequences in clinical settings. For example, an attacker could inject noise into the calcium imaging stream to cause the model to falsely predict low sleep pressure, potentially leading to unsafe prolonged wakefulness. To counter such threats, the inference pipeline includes an anomaly detector that monitors input distributions and flags deviations exceeding a threshold. If an anomaly is detected, the model falls back to a conservative default prediction or alerts the operator. These measures enhance the overall safety and trustworthiness of the system.

## **6. Policy and Governance Implications**

The deployment of deep learning models that predict neural ionic homeostasis changes raises significant policy and governance questions. At the most fundamental level, the data used to train and operate these models are highly sensitive. Transgenic animal data may be proprietary, and human neural data, when the technology eventually moves to clinical trials, would be considered protected health information under regulations such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in Europe. The system architecture must therefore incorporate privacy-preserving mechanisms. Federated learning, where model updates are shared without transmitting raw data, is one approach that can reduce privacy risks [23]. However, federated learning is challenging for time-series data with heterogeneous sampling rates and requires careful synchronization of gradient updates.

Another governance issue is algorithmic transparency. Regulators and clinicians need to understand how a deep learning model arrives at its predictions, especially when those predictions inform medical decisions. The inherent opacity of deep networks makes this difficult, but explainability techniques such as integrated gradients and attention visualization can provide partial insights [24]. The system should log the most influential input features for each prediction and present them in a human-readable format. Additionally, periodic validation studies should be conducted to compare model predictions against ground-truth

measurements under diverse conditions, and the results should be published to build community trust.

Intellectual property and open science also intersect in this domain. The development of genetically encoded biosensors and deep learning models involves contributions from multiple academic institutions and potentially commercial entities. Clear data-sharing agreements and licensing terms must be established to enable collaboration while protecting proprietary interests. Open-access repositories for both raw data and trained models can accelerate progress, but they also create risks of misuse, such as using the models to infer mental states without consent. Governance frameworks should include ethical review boards with neuroscience and AI expertise to oversee data access and use.

Policy considerations extend to the regulation of AI-based medical devices. In the United States, the Food and Drug Administration has issued guidance on software as a medical device, which would apply to predictive models that guide sleep management [25]. The proposed system would likely be classified as a Class II or III device, requiring premarket clearance or approval with evidence of safety and effectiveness. The model's robustness to domain shifts and adversarial inputs, as discussed in the previous section, becomes a regulatory requirement rather than an optional feature. Furthermore, post-market surveillance mechanisms must be implemented to monitor the model's performance in real-world settings and detect deterioration over time.

Finally, the broader societal implications of real-time sleep pressure prediction merit careful consideration. While the technology could improve sleep hygiene and reduce accidents caused by drowsiness, it also raises the specter of employer or insurer use to monitor employee sleep states without consent. Strong anti-discrimination laws and worker protections are essential to prevent misuse. Public dialogue involving neuroscientists, ethicists, policymakers, and affected communities should precede widespread deployment to ensure that the benefits of predictive deep learning for sleep health are distributed equitably.

## **7. Conclusion**

This paper has presented a comprehensive system-level framework for using deep learning to predict neural ionic homeostasis changes during sleep pressure accumulation. The proposed architecture integrates multimodal data acquisition, scalable cloud-edge processing, and sophisticated encoder-decoder models capable of forecasting ion concentration trajectories over timescales relevant to sleep regulation. Key design trade-offs between accuracy, latency, and energy efficiency are addressed through model compression, domain adaptation, and energy-aware scheduling. Robustness and fairness are enhanced through stratified sampling, adversarial debiasing, uncertainty quantification, and anomaly detection. Policy and governance considerations highlight the need for privacy-preserving data handling, algorithmic transparency, regulatory compliance, and ethical oversight.

The success of this framework depends on continued advances in biosensor technology, which will provide richer and more direct measurements of ionic dynamics. The genetically encoded ionic-stress sensor described in recent work [12] exemplifies the kind of tool needed to generate high-quality training data. Simultaneously, progress in efficient deep learning architectures will enable deployment on low-power devices suitable for wearable or implantable use. Large-scale collaborative efforts that pool data from multiple laboratories will be essential to train models with sufficient generalizability. By bridging molecular neurobiology and deep learning systems engineering, this work opens new avenues for real-

time sleep pressure monitoring, personalized sleep intervention, and a deeper understanding of the ionic underpinnings of sleep.

## References

1. Porkka-Heiskanen, T., & Kalinchuk, A. V. (2011). Adenosine, energy metabolism and sleep homeostasis. *Sleep Medicine Reviews*, 15(2), 123-135.
2. Ding, F., O'Donnell, J., Xu, Q., Kang, N., Goldman, N., & Nedergaard, M. (2016). Changes in the composition of brain interstitial ions control the sleep-wake cycle. *Science*, 352(6285), 550-555.
3. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
5. Glaser, J. I., Benjamin, A. S., Chowdhury, R. H., Perich, M. G., Miller, L. E., & Kording, K. P. (2020). Machine learning for neural decoding. *eNeuro*, 7(4).
6. Sors, A., Bonnet, S., Mirek, S., Vercueil, L., & Payan, Y. (2018). A convolutional neural network for sleep stage scoring from raw single-channel EEG. *Biomedical Signal Processing and Control*, 42, 107-114.
7. Borbély, A. A. (1982). A two process model of sleep regulation. *Human Neurobiology*, 1(3), 195-204.
8. O'Donnell, J., Ding, F., & Nedergaard, M. (2015). Distinct functional states of astrocytes during sleep and wakefulness: Is norepinephrine the master regulator? *Current Sleep Medicine Reports*, 1, 1-8.
9. Tononi, G., & Cirelli, C. (2014). Sleep and the price of plasticity: From synaptic and cellular homeostasis to memory consolidation and integration. *Neuron*, 81(1), 12-34.
10. Yoshida, H., & Tominaga-Yoshino, K. (2021). Intracellular pH and its role in sleep-wake regulation. *Neuroscience Research*, 172, 1-7.
11. Kawashima, T., Takahashi, N., Yoshida, J., & Ikegaya, Y. (2022). Proton dynamics in brain cells during sleep-wake cycle. *Journal of Neuroscience*, 42(18), 3750-3761.
12. Ji, Z., Liu, J., Wang, B., Wei, S., Bian, Y., Zeng, W., ... & Ma, D. K. (2026). A genetically encoded ionic-stress sensor reveals protons as a sleep driver. *bioRxiv*, 2026-01.
13. Achermann, P., & Dijk, D. J. (2003). Mathematical models of sleep regulation. *Frontiers in Bioscience*, 8, s683-s693.
14. Phan, H., Andreotti, F., Cooray, N., Chen, O. Y., & De Vos, M. (2019). Seqsleepnet: End-to-end hierarchical recurrent neural network for sequence-to-sequence automatic sleep staging. *IEEE Transactions on Biomedical Engineering*, 66(5), 1350-1361.
15. Sridhar, N., Shoeb, A., Stephens, P., Dement, W., & Yan, L. (2019). Deep learning for automated sleep staging using smartphone accelerometer data. *npj Digital Medicine*, 2, 87.

16. Casadei, R., & Karmacharya, S. (2023). Convolutional neural networks for calcium imaging analysis: A review. *Frontiers in Computational Neuroscience*, 17, 1154941.
17. Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6), 407-420.
18. Rossant, C., Kadir, S. N., Goodman, D. F. M., Schulman, J., Hunter, M. L. D., Saleem, A. B., ... & Harris, K. D. (2016). Spike sorting for large, dense electrode arrays. *Nature Neuroscience*, 19(4), 634-641.
19. Han, S., Mao, H., & Dally, W. J. (2016). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *International Conference on Learning Representations*.
20. Harris, J. J., & Jolivet, R. (2020). Ion dynamics and the neural basis of sleep homeostasis. *Current Opinion in Neurobiology*, 63, 150-157.
21. Zhang, B. H., Lemoine, B., & Mitchell, M. (2018). Mitigating unwanted biases with adversarial learning. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 223-228.
22. Wang, M., & Deng, W. (2018). Deep visual domain adaptation: A survey. *Neurocomputing*, 312, 135-153.
23. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. *Artificial Intelligence and Statistics*, 1273-1282.
24. Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. *International Conference on Machine Learning*, 3319-3328.
25. U.S. Food and Drug Administration. (2023). Marketing submission recommendations for a predetermined change control plan for artificial intelligence/machine learning (AI/ML)-enabled device software functions. FDA Guidance Document.