INTELLIGENT INCIDENT HYPERTENSION PREDICTION IN OBSTRUCTIVE SLEEP APNEA

Omid Halimi Milani[1], *Ahmet Enis Cetin*[1], *Bharati Prasad*[2]

[1]Department of Electrical and Computer Engineering, University of Illinois Chicago, Chicago, IL, USA
[2]Department of Medicine, Division of Pulmonary, Critical Care,
Sleep and Allergy, University of Illinois at Chicago, Chicago, IL, USA

ABSTRACT

Obstructive sleep apnea (OSA) is a significant risk factor for hypertension, primarily due to intermittent hypoxia and sleep fragmentation. Predicting whether individuals with OSA will develop hypertension within five years remains a complex challenge. This study introduces a novel deep learning approach that integrates Discrete Cosine Transform (DCT)-based transfer learning to enhance prediction accuracy. We are the first to incorporate all polysomnography signals together for hypertension prediction, leveraging their collective information to improve model performance. Features were extracted from these signals and transformed into a 2D representation to utilize pre-trained 2D neural networks such as MobileNet, EfficientNet, and ResNet variants. To further improve feature learning, we introduced a DCT layer, which transforms input features into a frequency-based representation, preserving spectral information and improving robustness. The frequency-domain approach improves generalization in small datasets. By strategically placing the DCT layer at deeper truncation depths within EfficientNet, our model achieved a best area under the curve (AUC) of 72.88%, demonstrating the effectiveness of frequency-domain feature extraction and transfer learning in predicting hypertension risk in OSA patients over a five-year period.

1. INTRODUCTION

Obstructive sleep apnea (OSA) is a condition characterized by intermittent hypoxia and sleep fragmentation, which propagates hypertension via mechanisms such as sympathetic activation and inflammation. OSA also raises the likelihood of developing hypertension during nighttime [1, 2].

Notably, specific patterns of apnea, such as Rapid Eye Movement (REM) sleep-related OSA, could contribute to OSA-associated hypertension [3]. Additional sleep-related factors, like reduced slow-wave sleep (SWS) and short sleep duration, are associated with hypertension, independent of OSA [4, 5].

Ren and colleagues conducted a study examining the association between sleep duration, obstructive sleep apnea (OSA), and hypertension in a group of 7,107 OSA patients and 1,118 primary snorers. The findings from polysomnography indicated that individuals who slept for 5 to 6 hours had a 45% risk of developing hypertension, while those who slept fewer than 5 hours had an 80% increased risk, independent of other influencing factors [6].

Treating obstructive sleep apnea (OSA) has been shown to lower the risk of developing hypertension [7]. However, accurately predicting the onset of hypertension in individuals with OSA remains a challenge due to the complex pathogenesis of hypertension.

We developed a deep learning model that predicts hypertension up to five years after OSA diagnosis using transformed polysomnography signals and static clinical features (e.g., age, sex, BMI, blood pressure). By integrating signal and clinical data, our approach supports precision medicine for early hypertension risk assessment.

We hypothesized that a comprehensive approach involving the simultaneous input of time-series physiological signals measuring sleep (EEG), ventilatory impairment and hypoxia, and cardiac autonomic dysregulation (electrocardiogram and photoplethysmography-derived heart rate variability and pulse transit time) could preserve the temporal correlations between multiple physiological perturbations in OSA and provide a robust prediction of incident hypertension [8]. These signals were included in polysomnography, a widely available diagnostic test for OSA. Thus, we extracted multiple features from the polysomnography signals in the Sleep Heart Health Study (SHHS) participants with moderate to severe OSA [9, 10].

We applied artifact removal and bandpass filtering before feature extraction.

The proposed methodology underwent a rigorous evaluation through a 10-fold cross-validation approach to examine the model's generalizability. The results were subsequently summarized, comparing models and methods to cutting-edge approaches.

Our main contributions are summarized as follows:

- We develop a DCT-based convolution framework that replaces the complex-valued Discrete Fourier Trans-

form (DFT) with a real-valued, orthogonal DCT, thereby simplifying convolution operations and preserving crucial frequency information in polysomnography signals.

- We introduce threshold-based nonlinearities (soft and hard thresholding) in the DCT domain, preventing the loss of important negative-frequency coefficients that standard ReLU would discard.

- We embed the DCT layer within a truncated EfficientNet architecture, using time-windowed feature extraction and transfer learning to efficiently process multichannel polysomnography data.

- We demonstrate that this end-to-end model accurately predicts long-term hypertension risk, providing a robust framework for precision healthcare in OSA.

2. RELATED WORKS

In recent years, there has been a growing recognition of the limitations associated with conventional manual sleep stage scoring, which simplifies the analysis of electroencephalogram (EEG) temporospectral and frequency domains. This scoring method is inherently subjective and can lead to variations between different scorers due to the application of visual-based rules [11]. Researchers used power spectral density (PSD) analysis of EEG signals, which examines sleep EEG microarchitecture. This approach enables the decomposition of EEG brain waves across various power frequency bands, ranging from slow wave activity (delta EEG power, 1–4 Hz) to fast-frequency activity (beta EEG power, 18–30 Hz), achieved through fast Fourier transform algorithms. At a microarchitecture level, slow wave sleep (SWS) is characterized by high delta power, indicative of deep sleep. Quantitative EEG analysis may yield more sensitive biomarkers for adverse health outcomes in OSA compared to traditional sleep scoring methods [12, 13, 14, 15].

Low delta power during non-REM sleep has been linked to increased hypertension risk, reinforcing the role of SWS in blood pressure regulation [16]. Another study using only SpO_2 features with time-frequency analysis achieved 84.3% AUC in predicting hypertension among OSA patients [17].

Ruitong et al. adopted a pulmonary physiology-based approach to predicting the onset of hypertension by including pulmonary function measurements and polysomnography-derived indices using a penalized regression and Elastic Net model [18]. A recent study introduced cSP (sleep and pulmonary) phenotypes, which combines spirometry and overnight polysomnography measures to predict hypertension occurrence in the SHHS [12]. The SHHS dataset encompasses a variety of physiological signals, including sleep EEG, electrocardiogram (ECG), electromyogram (EMG), ventilatory effort, nasal airflow, photoplethysmography-derived oximetry, snoring, and body position. Employing rigorous signal processing techniques, such as filtering, segmentation, and feature extraction, our analysis aimed to unravel the intricate patterns embedded within these signals [9, 10].

3. METHODOLOGY

We propose a feature-based deep learning method for predicting incident hypertension in patients with obstructive sleep apnea (OSA). Convolution operations are typically performed in the time (or spatial) domain, but they can also be implemented more efficiently in the frequency domain via element-wise multiplication. Let $y = w * x$ be the convolution of a filter w and signal x. By applying the DFT, we obtain:

$$Y[k] = W[k] \cdot X[k], \tag{1}$$

where Y, W, and X are the DFTs of y, w, and x, respectively. However, the DFT produces complex-valued coefficients, complicating the use of conventional non-linearities such as ReLU. To avoid this, we use the Discrete Cosine Transform (DCT), a purely real and orthogonal transform, which preserves signal energy in the frequency domain while reducing redundancy.

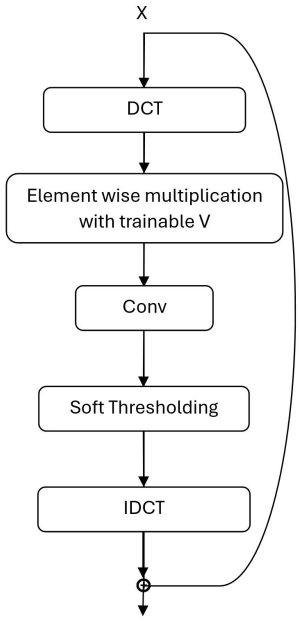

Fig. 1: The structure of the DCT block. V is a trainable matrix that adapts frequency features based on training data.

A key feature of our approach is the replacement of traditional convolutions with a DCT-based convolution block (illustrated in Fig. 1). Rather than applying ReLU activations in the DCT

domain—which would discard negative-frequency information—we incorporate a soft thresholding operator. Originally introduced in image denoising, soft thresholding effectively prunes smaller (potentially noisy) DCT coefficients while retaining both positive and negative high-amplitude components [19]. In our PyTorch implementation, it is defined as:

$$\mathrm{SoftThreshold}(x, \tau) = \mathrm{sign}(x)\ \max\big(|x| - \tau, 0\big). \quad (2)$$

By applying soft thresholding in place of ReLU, we preserve crucial frequency details that might otherwise be lost, thereby improving the expressive capacity of the model.

Another important characteristic of the DCT layer is its orthogonality. Because each transform basis (cosine wave) is orthogonal to the others, the resulting representation can decorrelate input features, often leading to improved generalization and numerical stability compared to complex-valued transforms.

To reduce computational overhead, we do not operate on raw polysomnography signals directly. Instead, each signal (EEG, ECG, respiratory, etc.) is first preprocessed (e.g., 0.5–40 Hz for EEG) and then segmented into time windows. From each window, we extract a set of representative features, including:

- Counts and durations of respiratory events (apneas, hypopneas) and arousals,
- Statistical descriptors (mean, standard deviation, skewness, kurtosis),
- Heart rate variability (HRV) metrics derived from ECG signals.

Arranging the extracted features from windowed segments sequentially preserves coarse temporal structure. Subsequent convolutional layers learn temporal-frequency patterns across windows, which are further enhanced by the DCT layer.

HRV Feature Computation. After detecting R-peaks within each window, we obtain the sequence of R-R intervals, $\{RR_i\}_{i=1}^{N}$. Let $\overline{RR}$ be the mean R-R interval:

$$\overline{RR} = \frac{1}{N}\sum_{i=1}^{N} RR_i. \quad (3)$$

We then compute the following standard HRV metrics:

$$\mathrm{SDNN} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\big(RR_i - \overline{RR}\big)^2}, \quad (4)$$

$$\mathrm{RMSSD} = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N-1}\big(RR_{i+1} - RR_i\big)^2}, \quad (5)$$

$$\mathrm{pNN50} = \frac{\mathrm{NN50}}{N-1} \times 100\%, \quad (6)$$

$$\mathrm{NN50} = \sum_{i=1}^{N-1} \mathbf{1}\big(|RR_{i+1} - RR_i| > 50\ \mathrm{ms}\big). \quad (7)$$

Here, N is the number of valid R-R intervals within the window. SDNN measures overall variability, RMSSD captures short-term fluctuations, and pNN50 reflects the percentage of adjacent R-R intervals that differ by more than 50 ms.

We concatenate all window-level features into a $2D$ array (features $\times$ windows), treating it as a pseudo-image [20]. This 2D format enables use of CNNs with DCT-enhanced features.

By replacing DFT with real-valued DCT and using soft thresholding, our method preserves informative frequency components while leveraging a truncated CNN and spectral features for effective hypertension prediction in OSA patients.

4. RESULTS

To evaluate the performance of our *feature-based* approach, we employed 10-fold cross-validation with various time window lengths ranging from 9 to 60 minutes. We present here the results for 60-minute and 10-minute intervals, as well as a brief investigation into optimized shorter windows around 10 minutes. We then discuss how adding a Discrete Cosine Transform (DCT) block at different depths within a truncated EfficientNet affects performance.

4.1. 60-Minute Interval Results

Table 1 summarizes the accuracy and AUC for the 60-minute interval. Among the feature-based models, Feature_EffNet-B0 achieved the highest accuracy (68.66%) and AUC (68.66%). Other models had lower sensitivity but higher specificity

Table 1: Performance of Models at the 60-Minute Interval

Model	Accuracy (%)	AUC (%)
Feature_EffNet-B0	**68.66**	**68.66**
Feature_ResNet-10	62.69	59.70
Feature_ResNet-18	61.19	55.22
Feature_MobileNet-v2	65.67	61.19

4.2. 10-Minute Interval Results

A shorter 10-minute window offers more granular insight into sleep data. As shown in Table 2, Feature_EffNet-B0 achieved the highest accuracy (68.66%) and AUC (71.64%).

Incorporating static features similarly provided small boosts in certain metrics. 10-minute windows offer a good trade-off between data and performance.

Table 2: Performance of Models at the 10-Minute Interval

Model	Accuracy (%)	AUC (%)
Feature_EffNet-B0	**68.66**	**71.64**
Feature_ResNet-10	62.69	60.70
Feature_ResNet-18	61.19	57.31
Feature_MobileNet-v2	65.67	63.43

4.3. DCT Block Placement at Different Truncation Depths

We evaluated how inserting the DCT2D layer at different truncation depths in EfficientNet-B0 affects predictive performance. Specifically, we integrated the DCT block after the third, fourth, fifth, or sixth building block in the network. Table 3 presents the best observed accuracy and AUC for each placement.

Table 3: Best Accuracy and AUC by Inserting the DCT2D Block at Different EfficientNet-B0 Depths

Model	Accuracy (%)	AUC (%)
DCT@3 (after 3rd block)	68.66	70.81
DCT@4 (after 4th block)	69.66	70.73
DCT@5 (after 5th block)	68.12	72.79
DCT@6 (after 6th block)	69.86	72.88

As shown in Table 3, placing the DCT layer at deeper levels (DCT@5 or DCT@6) yields higher AUC values (72.29% and 72.88%, respectively), with DCT@6 also achieving the top accuracy (69.86%). This suggests that mid- to late-stage feature maps in EfficientNet-B0 may benefit more from the frequency-domain transform, likely due to increasingly abstract representations of the polysomnography data at deeper layers. Ultimately, the choice between prioritizing accuracy or maximizing AUC may depend on clinical considerations, such as avoiding false negatives versus improving overall predictive discrimination.

4.4. Comparison with State-of-the-Art Methods

Table 4 compares our best-performing DCT-based approach to established methods, including cSPPSG and AHI [18]. Incorporating the DCT layer at depth 5 or 6 outperforms these baselines in terms of AUC, underscoring the advantage of combining frequency-domain transformations with threshold-based nonlinearities for improved hypertension-risk prediction in OSA.

In conclusion, integrating the DCT layer at deeper truncation depths within EfficientNet-B0 yields meaningful gains in AUC while maintaining high accuracy. This finding highlights the synergy between advanced CNN architectures and frequency-domain transformations for predicting long-term hypertension in OSA patients.

Table 4: Comparison with State-of-the-Art Models

Model	Accuracy (%)	AUC (%)
cSPPSG [18]	-	71
AHI [18]	-	67
Ours	69.86	**72.88**

5. CONCLUSION

This study introduced a DCT-enhanced deep learning framework to predict incident hypertension in obstructive sleep apnea (OSA) patients using polysomnography data. By embedding threshold-based nonlinearities in the DCT domain within a truncated EfficientNet backbone, we preserved essential negative-frequency information and effectively leveraged multi-signal features. Our best-performing configuration achieved 69.86% accuracy and a 72.88% AUC, surpassing existing models such as cSPPSG and AHI, which achieved 71% and 67% AUC respectively. These results highlight the potential of combining spectral-domain methods with CNNs for hypertension prediction.

Limited data may affect generalizability. Future work will focus on expanding data availability, and further improving the model architecture. By addressing these areas, we aim to enhance predictive performance and support earlier, more precise interventions for OSA patients at risk of developing hypertension.

6. REFERENCES

[1] Paul E Peppard, Terry Young, Mari Palta, and James Skatrud, "Prospective study of the association between sleep-disordered breathing and hypertension," *New England Journal of Medicine*, vol. 342, no. 19, pp. 1378–1384, 2000.

[2] T Dudenbostel and DA Calhoun, "Resistant hypertension, obstructive sleep apnoea and aldosterone," *Journal of human hypertension*, vol. 26, no. 5, pp. 281–287, 2012.

[3] Babak Mokhlesi, Laurel A Finn, Erika W Hagen, Terry Young, Khin Mae Hla, Eve Van Cauter, and Paul E Peppard, "Obstructive sleep apnea during rem sleep and hypertension. results of the wisconsin sleep cohort," *American journal of respiratory and critical care medicine*, vol. 190, no. 10, pp. 1158–1167, 2014.

[4] Sogol Javaheri, Ying Y Zhao, Naresh M Punjabi, Stuart F Quan, Daniel J Gottlieb, and Susan Redline,

"Slow-wave sleep is associated with incident hypertension: the sleep heart health study," *Sleep*, vol. 41, no. 1, pp. zsx179, 2018.

[5] Luyao Wang, Youdan Hu, Xin Wang, Shu Yang, Weizhong Chen, and Ziqian Zeng, "The association between sleep duration and hypertension: a meta and study sequential analysis," *Journal of Human Hypertension*, vol. 35, no. 7, pp. 621–626, 2021.

[6] Rong Ren, Naima Covassin, Linghui Yang, Yun Li, Ye Zhang, Junying Zhou, Lu Tan, Taomei Li, Xiao Li, Yanyan Wang, et al., "Objective but not subjective short sleep duration is associated with hypertension in obstructive sleep apnea," *Hypertension*, vol. 72, no. 3, pp. 610–617, 2018.

[7] José M Marin, Alvar Agusti, Isabel Villar, Marta Forner, David Nieto, Santiago J Carrizo, Ferran Barbé, Eugenio Vicente, Ying Wei, F Javier Nieto, et al., "Association between treated and untreated obstructive sleep apnea and risk of hypertension," *Jama*, vol. 307, no. 20, pp. 2169–2176, 2012.

[8] Shitao Zheng and Dongrui Wu, "Semi-supervised domain adaptation for eeg-based sleep stage classification," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 1776–1780.

[9] Stuart F Quan, Barbara V Howard, Conrad Iber, James P Kiley, F Javier Nieto, George T O'Connor, David M Rapoport, Susan Redline, John Robbins, Jonathan M Samet, et al., "The sleep heart health study: design, rationale, and methods," *Sleep*, vol. 20, no. 12, pp. 1077–1085, 1997.

[10] Guo-Qiang Zhang, Licong Cui, Remo Mueller, Shiqiang Tao, Matthew Kim, Michael Rueschman, Sara Mariani, Daniel Mobley, and Susan Redline, "The national sleep research resource: towards a sleep data commons," *Journal of the American Medical Informatics Association*, vol. 25, no. 10, pp. 1351–1358, 2018.

[11] Hartmut Schulz, "Rethinking sleep analysis: comment on the aasm manual for the scoring of sleep and associated events," *Journal of Clinical Sleep Medicine*, vol. 4, no. 2, pp. 99–103, 2008.

[12] Sarah L Appleton, Andrew Vakulin, Angela D'Rozario, Andrew D Vincent, Alison Teare, Sean A Martin, Gary A Wittert, R Doug McEvoy, Peter G Catcheside, and Robert J Adams, "Quantitative electroencephalography measures in rapid eye movement and nonrapid eye movement sleep are associated with apnea–hypopnea index and nocturnal hypoxemia in men," *Sleep*, vol. 42, no. 7, pp. zsz092, 2019.

[13] Angela L D'Rozario, Nathan E Cross, Andrew Vakulin, Delwyn J Bartlett, Keith KH Wong, David Wang, and Ronald R Grunstein, "Quantitative electroencephalogram measures in adult obstructive sleep apnea–potential biomarkers of neurobehavioural functioning," *Sleep medicine reviews*, vol. 36, pp. 29–42, 2017.

[14] Bastien Lechat, Kristy L Hansen, Yohannes Adama Melaku, Andrew Vakulin, Gorica Micic, Robert J Adams, Sarah Appleton, Danny J Eckert, Peter Catcheside, and Branko Zajamsek, "A novel electroencephalogram-derived measure of disrupted delta wave activity during sleep predicts all-cause mortality risk," *Annals of the American Thoracic Society*, vol. 19, no. 4, pp. 649–658, 2022.

[15] Jesse L Parker, Yohannes Adama Melaku, Angela L D'Rozario, Gary A Wittert, Sean A Martin, Peter G Catcheside, Bastien Lechat, Alison J Teare, Robert J Adams, Sarah L Appleton, et al., "The association between obstructive sleep apnea and sleep spindles in middle-aged and older men: a community-based cohort study," *Sleep*, vol. 45, no. 3, pp. zsab282, 2022.

[16] Mathieu Berger, Andrew Vakulin, Camila Hirotsu, Nicola Andrea Marchi, Geoffroy Solelhac, Virginie Bayon, Francesca Siclari, José Haba-Rubio, Julien Vaucher, Peter Vollenweider, et al., "Association between sleep microstructure and incident hypertension in a population-based sample: The hypnolaus study," *Journal of the American Heart Association*, vol. 11, no. 14, pp. e025828, 2022.

[17] Jingyuan You, Juan Li, Xiaoyu Li, Haojie Li, Jinying Tu, Yuhuan Zhang, Jiandong Gao, Ji Wu, and Jingying Ye, "Risk-prediction model for incident hypertension in patients with obstructive sleep apnea based on spo2 signals," in *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, 2023, pp. 1–4.

[18] Ruitong Li, Michael Rueschman, Daniel J Gottlieb, Susan Redline, and Tamar Sofer, "A composite sleep and pulmonary phenotype predicting hypertension," *EBioMedicine*, vol. 68, 2021.

[19] David L Donoho, "De-noising by soft-thresholding," *IEEE transactions on information theory*, vol. 41, no. 3, pp. 613–627, 2002.

[20] Omid Halimi Milani, Tu Nguyen, Ankit Parekh, Ahmet Enis Cetin, and Bharati Prasad, "0537 incident hypertension prediction in obstructive sleep apnea using machine learning," *Sleep*, vol. 46, no. Supplement_1, pp. A236–A237, 2023.