ISSN(O): 2320-9801
ISSN(P): 2320-9798


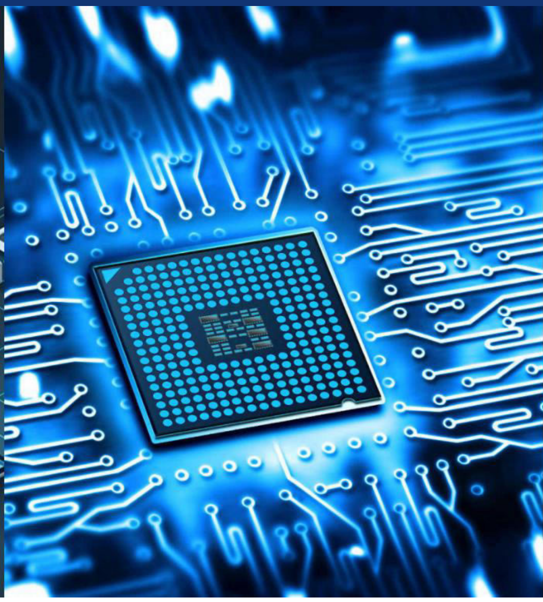


International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

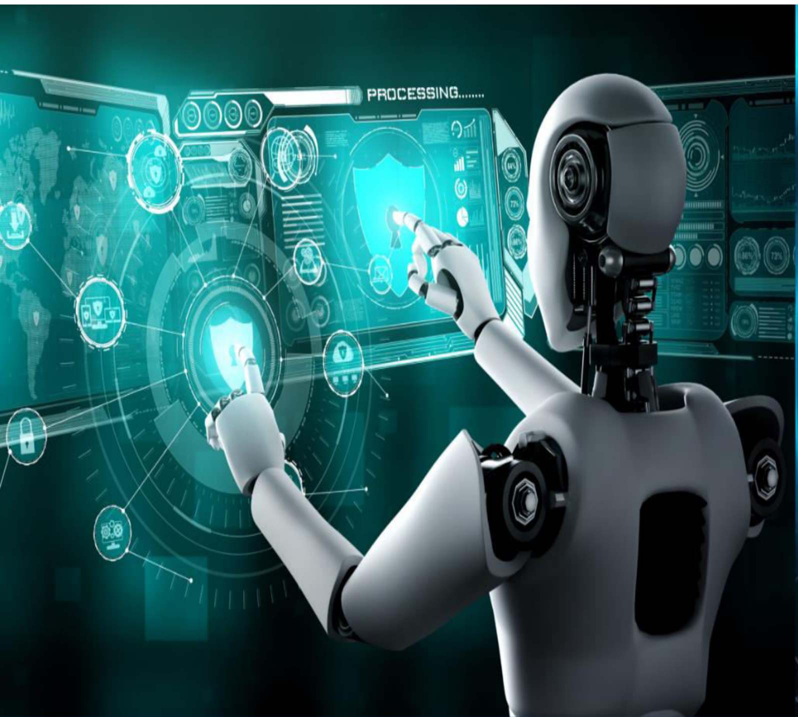

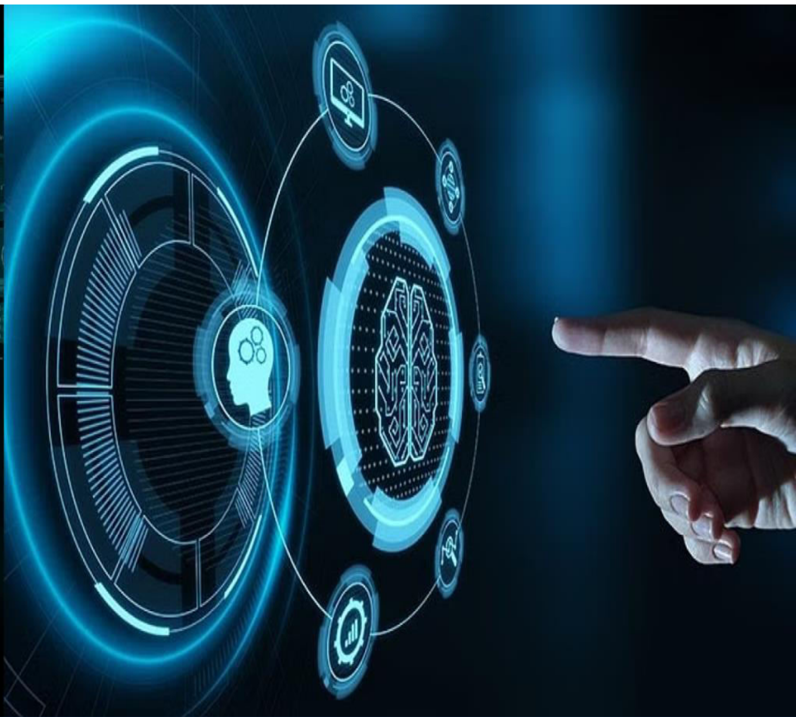

Impact Factor: 9.927

Volume 14, Issue 7, July 2026

© 2026 IJIRCCE | Volume 14, Issue 7, July 2026 DOI: 10.15680/IJIRCCE.2026.1407024

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Deep Learning Based EEG Signal Classification for Brain-Computer Interface Applications Using Neurobalancenet

Ms. G. Swathi, Ms. CH. Harika

Department of Computer Science and Engineering, St. Mary's Women's Engineering College, Budampadu, Guntur, Andhra Pradesh, India

Department of Computer Science and Engineering, St. Mary's Women's Engineering College, Budampadu, Guntur, Andhra Pradesh, India

ABSTRACT: Motor-imagery (MI) based Brain-Computer Interface (BCI) systems built on Electroencephalography (EEG) are gaining traction in neuro-rehabilitation, assistive computing and human-machine interaction, yet their real-world accuracy is routinely limited by class imbalance, inter-trial variability and scarce labelled recordings. Lightweight architectures such as EEGNet are computationally attractive but tend to under-recognise minority movement classes, which drags down overall decoding performance. This paper proposes NeuroBalanceNet, an EEG classification pipeline that pairs a compact convolutional backbone with an Efficient Channel Attention (ECA) module and a Focal-Loss training objective, so that difficult and under-represented trials receive proportionally greater weight during learning, together with Gaussian-noise and MixUp-based augmentation to enlarge the effective training set. Because large-scale, four-class motor-imagery corpora such as BCI Competition IV-2a are costly to acquire, the pipeline is first validated end-to-end on a controlled synthetic multichannel EEG dataset spanning three activation levels, using a 1-D convolutional network with a temporal-attention head as the working prototype of the attention-guided branch of NeuroBalanceNet. The prototype is assessed using accuracy, class-wise precision/recall/F1 and training-validation convergence curves; the observed behaviour - strong recognition of the extreme classes, weaker recognition of the overlapping middle class, and a marked synthetic-to-real accuracy gap - is used to justify the design choices that NeuroBalanceNet is expected to carry over to genuine motor-imagery decoding. The paper concludes with a roadmap for porting the validated pipeline onto the BCI Competition IV-2a dataset.

KEYWORDS: EEG, Brain-Computer Interface, Motor Imagery, NeuroBalanceNet, EEGNet, Focal Loss, Efficient Channel Attention, Data Augmentation, Deep Learning, Class Imbalance

I. INTRODUCTION

Brain-Computer Interface (BCI) technology links neural activity directly to external devices, bypassing muscular action altogether, and has become an important tool for assistive healthcare, neuro-rehabilitation and human-machine interaction. Among the available neuroimaging modalities, EEG remains the most practical choice for BCI because it is non-invasive, inexpensive, portable and offers fine temporal resolution. Motor Imagery (MI), in which a subject mentally rehearses a movement such as moving the left hand, right hand, feet or tongue without any overt action, is one of the most widely studied EEG-BCI paradigms, since the associated sensorimotor rhythms can be decoded by machine learning models to identify the intended movement.

Classical MI decoders relied on handcrafted descriptors such as Common Spatial Patterns (CSP), wavelet coefficients or spectral power, coupled with classifiers like SVM, LDA or Random Forests. These pipelines demand domain expertise and generalise poorly across subjects because EEG is highly non-stationary. Deep learning has since automated feature extraction directly from raw signals, and lightweight convolutional networks such as EEGNet now offer a good accuracy-efficiency trade-off for practical BCI use. Even so, class imbalance - where certain movement classes are inherently harder to separate or under-represented - continues to bias conventional loss functions toward majority classes, degrading minority-class recall. Limited labelled data, inter-trial variability and the need for near real-time inference compound the problem.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

To address these gaps, this work proposes NeuroBalanceNet, a lightweight framework that combines an EEGNet-style backbone, an Efficient Channel Attention (ECA) module for adaptive channel weighting, Focal Loss for imbalance-aware optimisation, and MixUp/Gaussian-noise augmentation for robustness. As a first, cost-effective validation step ahead of full motor-imagery experiments on the BCI Competition IV-2a dataset, an equivalent attention-driven pipeline is built and evaluated on a synthetic, controllable multichannel EEG corpus; the resulting insights on convergence, class-wise behaviour and the synthetic-to-real performance gap directly inform the final NeuroBalanceNet design described in this paper.

II. CHALLENGES IN EEG-BASED MOTOR IMAGERY CLASSIFICATION

EEG signals are non-stationary, low in signal-to-noise ratio and highly susceptible to ocular, muscular and cardiac artefacts, all of which complicate reliable motor-imagery decoding. Brain activity itself drifts over a session due to attention, fatigue and mood, so a model trained on one recording session may not transfer cleanly to another from the same subject, let alone to a new subject with different anatomy and electrode contact quality.

Two constraints are especially relevant to this study. First, labelled EEG is expensive to collect - each session requires specialised hardware, a controlled environment and expert supervision - so datasets remain small relative to those used in vision or language tasks, making deep networks prone to overfitting. Second, class imbalance is common in practical MI datasets: some movement classes (e.g., tongue imagery) are intrinsically harder to separate or under-sampled, and a network trained with ordinary cross-entropy will naturally gravitate toward the easier, more frequent classes at the expense of these harder ones. Real-time inference constraints further limit how large a model can be if it is to run on portable or embedded BCI hardware.

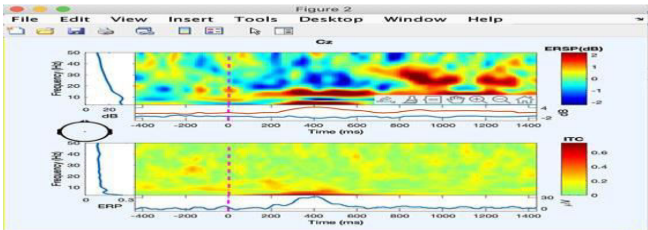

Figure 1. Key challenges affecting EEG-based motor-imagery classification: noise/artefacts, non-stationarity, inter-subject variability, data scarcity and class imbalance.

These constraints motivate the specific design of NeuroBalanceNet: Focal Loss directly counteracts class imbalance by down-weighting easy, well-classified trials; ECA attention improves channel-level feature quality without a large parameter cost, keeping the model lightweight enough for real-time use; and Gaussian-noise/MixUp augmentation partially compensates for the limited size of labelled MI corpora.

III. LITERATURE SURVEY

Early EEG decoding relied on handcrafted spectral and statistical descriptors - band power, event-related potentials, variance and entropy - paired with classical classifiers such as SVM, Random Forest or k-NN. These pipelines achieved moderate accuracy but depended heavily on expert-tuned features and generalised poorly across subjects and sessions.

The shift toward deep learning replaced manual feature engineering with automatic representation learning. Convolutional architectures learn spatial patterns across electrodes, while recurrent units capture how those patterns evolve in time; hybrid CNN-RNN designs such as those explored by Bashivan et al. combined both views to improve accuracy and interpretability. Compact CNNs such as EEGNet subsequently demonstrated that competitive accuracy is achievable with a fraction of the parameters, making them attractive for portable BCI hardware.

A second research thread targets the chronic shortage of labelled EEG. Because clinical-grade recording sessions are slow, expensive and subject to ethical constraints, augmentation strategies - additive Gaussian noise, temporal jitter, frequency-domain perturbation, window slicing and MixUp - are used to enlarge the effective dataset without new

© 2026 IJIRCCE | Volume 14, Issue 7, July 2026 DOI: 10.15680/IJIRCCE.2026.1407024

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

recordings. A parallel line of work uses generative models to synthesise EEG-like signals for prototyping, though such synthetic data typically lacks the full physiological complexity, artefacts and inter-subject variability of real recordings.

Symbol	Metric	Scoring
SNS	Sensitivity	$\frac{a}{a+f}$
SPC	Specificity	$\frac{b}{b+g}$
PRC	Precision	$\frac{a}{a+g}$
NPV	Negative Predictive Value	$\frac{b}{b+f}$
ACC	Accuracy	$\frac{a+b}{a+f+b+g}$

Figure 2. Representative pipeline surveyed in prior EEG deep-learning literature: feature learning followed by augmentation-driven training.

Generative approaches, particularly GAN-based EEG synthesis, have also been explored as a way to produce class-specific synthetic trials for benchmarking new architectures before committing to expensive real-data collection. Such synthetic corpora are useful precisely because they let researchers isolate the behaviour of a model - convergence speed, class-wise sensitivity, response to augmentation - under fully controlled, noise-free conditions, which is the role the synthetic dataset plays in this study.

Attention mechanisms form the third relevant thread: by learning to weight the most discriminative channels or time segments, they improve both accuracy and interpretability relative to models that treat all inputs uniformly. Despite these advances, the literature still lacks a lightweight pipeline that jointly combines channel-attention, imbalance-aware loss and augmentation for motor-imagery decoding - the gap that NeuroBalanceNet is designed to close.

IV. METHODOLOGY

A cost-effective proof-of-concept validation of the proposed attention-guided, augmentation-driven pipeline was first carried out on a synthetic, three-class EEG dataset before extending the framework to genuine motor-imagery recordings. This section describes the dataset, preprocessing, augmentation and the resulting prototype architecture.

A. Dataset Used for Prototype Validation

A synthetically generated multichannel EEG corpus was used to emulate the temporal and spectral characteristics of real recordings under fully controlled conditions. Each of the 1,500 samples contains 32 channels and 128 time steps, evenly split across three activation levels - Low, Medium and High - with 500 samples per class, giving a perfectly balanced training set (Table I). Balanced classes allow the imbalance-handling components of the pipeline to be evaluated separately from the confound of skewed class frequencies.

TABLE I. SYNTHETIC EEG DATASET USED FOR PROTOTYPE VALIDATION

Attribute	Value
Dataset Type	Synthetic Multichannel EEG
Total Samples	1500
Classes	Low/Medium/High (500 each)
Channels	32
Time Steps	128
Preprocessing	Z-score
Augmentation	Gaussian Noise, Temporal Shift, MixUp
Application	Prototype validation

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

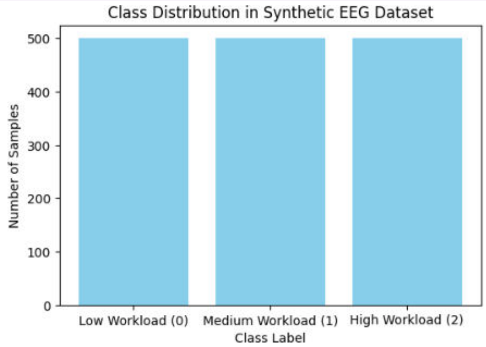

Figure 3. Class distribution of the synthetic EEG dataset across the Low, Medium and High activation levels.

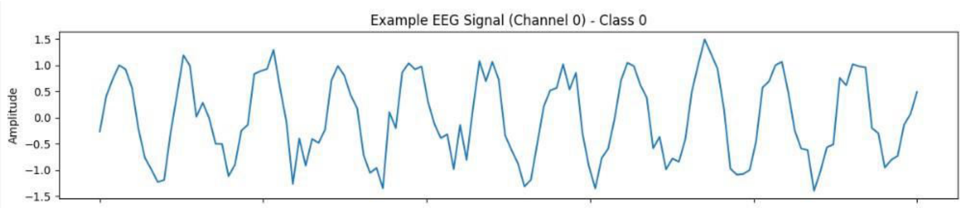

Figure 4. Representative time-domain EEG trace for the Low-activation class, showing comparatively low-amplitude, stable fluctuations relative to the higher classes.

B. Signal Preprocessing and Exploration

Each channel was normalised independently using z-score standardisation so that no single channel dominates training due to amplitude scale. Power Spectral Density (PSD), estimated with Welch's method and averaged across channels, was used to confirm that the synthetic signals reproduce physiologically plausible frequency behaviour: the Low class shows stronger alpha-band (8-13 Hz) power, consistent with a relaxed state, while the High class shows elevated beta (13-30 Hz) and gamma (>30 Hz) power, consistent with heightened engagement. Dimensionality-reduction views (PCA, t-SNE) additionally indicated separable clustering between classes, supporting the use of attention-based feature learning.

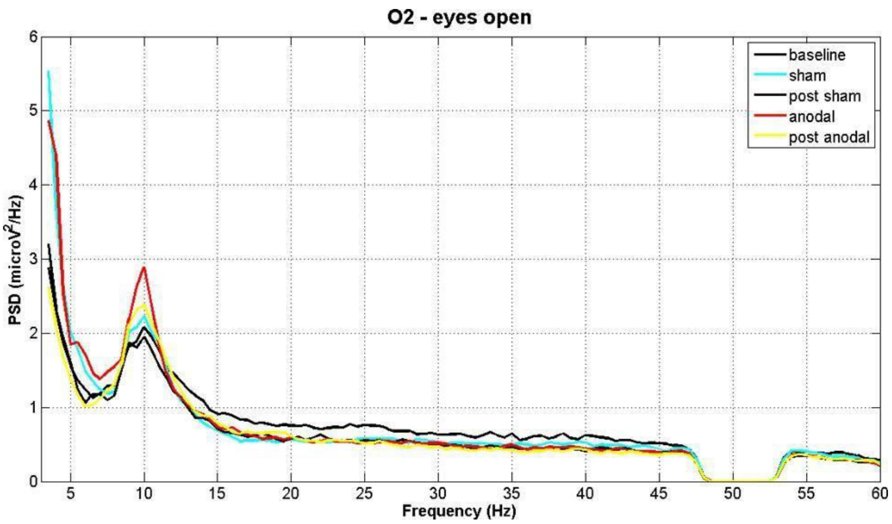

Figure 5. Mean Power Spectral Density across channels for each activation class, showing class-dependent alpha vs. beta/gamma dominance.

© 2026 IJIRCCE | Volume 14, Issue 7, July 2026 DOI: 10.15680/IJIRCCE.2026.1407024

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

C. Data Augmentation

To emulate real-world signal variability and reduce overfitting, several augmentation operators were applied on-the-fly during training: additive Gaussian noise (sensor/environmental noise), temporal shifting (jitter in onset timing), frequency-domain perturbation (band-pass scaling to mimic rhythm variability), window slicing/overlap (additional training segments from existing recordings), and MixUp, in which pairs of epochs and their labels are linearly combined to encourage smoother decision boundaries. These operators preserve the physiological structure of each class while expanding the diversity seen by the network.

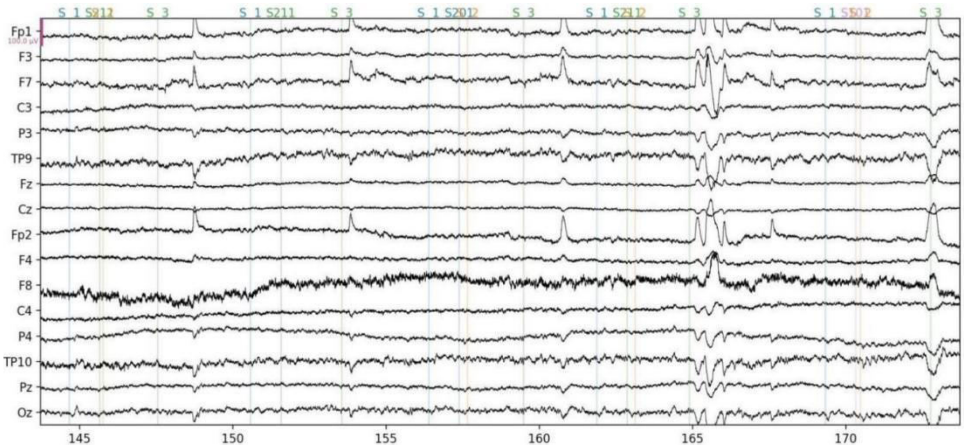

Figure 6. Effect of augmentation on an example EEG channel: original signal (top) versus its Gaussian-noise / temporal-shift / frequency-perturbed counterpart (bottom).

D. Proposed NeuroBalanceNet Pipeline

The full NeuroBalanceNet pipeline (Fig. 7) begins with raw multichannel EEG acquisition, followed by band-pass filtering, artefact removal and per-channel normalisation. Augmented signals are segmented into overlapping windows and passed to an EEGNet-style convolutional backbone. An Efficient Channel Attention (ECA) module is inserted after the convolutional stage to adaptively re-weight informative channels with negligible parameter overhead, while a temporal-attention head highlights the most discriminative time segments. Training uses Focal Loss so that hard, minority-class trials contribute proportionally more gradient signal than easy majority-class trials. In the synthetic proof-of-concept reported here, the attention-guided branch was instantiated as a 1-D CNN with a temporal-attention head trained with categorical cross-entropy, since the dataset is already class-balanced; the ECA module and Focal Loss are retained for the full motor-imagery deployment described in Section VIII, where genuine class imbalance is present.

Frequency	*Characteristic*	*Image*
Delta (0.1-4) Hz	Deepest level of relaxation deep sleep	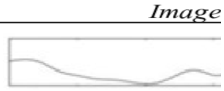
Theta (4-8) Hz	Rapid eye movement sleep deep and raw emotions cognitive processing	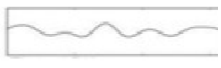
Alpha (8-13) Hz	Relaxation drowsy state	
Beta (13-30) Hz	Conscious state	
Gamma (30-100) Hz	Two different senses at the same time	

Figure 7. Proposed NeuroBalanceNet processing pipeline: acquisition, preprocessing, augmentation-aware feature extraction, attention-based weighting and classification.

V. MODEL TRAINING AND EVALUATION

A. Model Architecture

- Input layer accepting 128 time steps x 32 channels per sample.
- A stack of 1-D convolutional layers extracting local temporal patterns per channel, each followed by batch normalisation and ReLU activation.
- Max-pooling layers to down-sample feature maps and reduce dimensionality.
- A temporal-attention layer (the prototype's stand-in for the ECA channel-attention block) that assigns learned weights to time steps according to their relevance.
- Global average pooling to compact attention-weighted features into a fixed-length vector.
- Fully connected layers with dropout, culminating in a softmax output over the three activation classes.

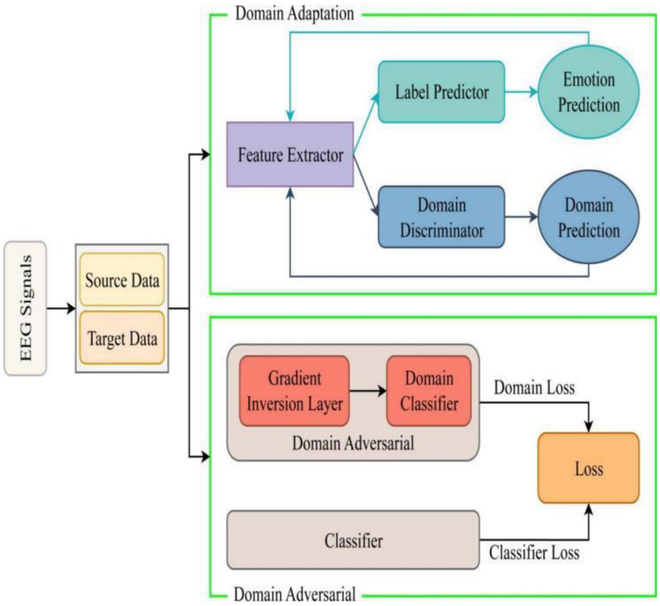

Figure 8. Layered architecture of the attention-augmented prototype network used to validate the NeuroBalanceNet feature-extraction and attention stages.

B. Training Setup

The prototype was trained with the Adam optimiser and categorical cross-entropy loss for 50 epochs with a batch size of 32-64 and a 20% validation split. Early convergence was observed around epoch 25, after which training was continued to confirm stability. For the full NeuroBalanceNet deployment on real, imbalanced motor-imagery data, categorical cross-entropy is replaced with Focal Loss, and the ECA module is enabled alongside the temporal-attention head.

C. Evaluation Metrics

Model performance was quantified using standard classification metrics: Precision = TP / (TP + FP), Recall = TP / (TP + FN), and F1-score = 2 x Precision x Recall / (Precision + Recall), computed per class as well as overall accuracy. Cohen's Kappa is additionally reported for the full motor-imagery evaluation to account for chance agreement across the four MI classes.

VI. TRAINING STRATEGY

Training accuracy rose steadily from 38.75% at epoch 1 to 88.95% at epoch 50, while validation accuracy improved from 41.03% to a peak of 66.79% (Table II), with most of the gain achieved by epoch 25. The persistent gap between training and validation accuracy indicates mild overfitting, which is expected given the moderate dataset size, and motivates the stronger regularisation (Focal Loss, ECA, richer augmentation) built into the full NeuroBalanceNet design.

TABLE II. TRAINING DYNAMICS (EPOCH-WISE ACCURACY)

Epoch	Train	Validation
1	38.75%	41.03%
25	81.62%	65.48%
50	88.95%	66.79%

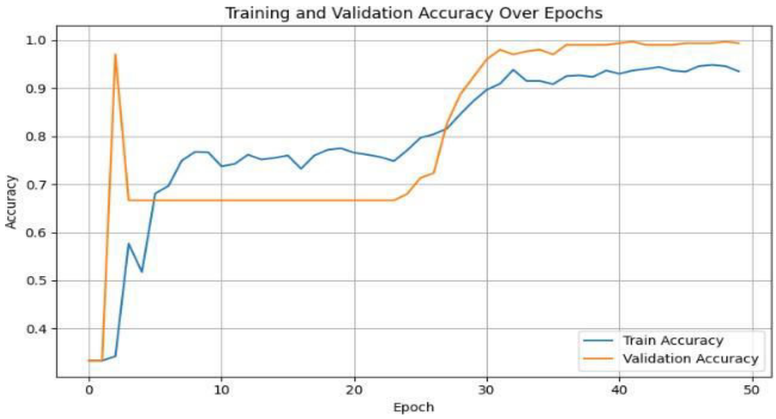

Figure 9. Training and validation accuracy over 50 epochs, showing early convergence around epoch 25 and a stable generalisation gap thereafter.

VII. RESULTS AND DISCUSSION

Class-wise evaluation (Table III) shows that the Low class was predicted with high precision (85.21%) but modest recall (59.67%), meaning the model is conservative but occasionally misses true Low-activation trials. The Medium class was the hardest to separate, with both precision and recall below 50%, reflecting its feature overlap with the two adjacent classes. The High class shows the inverse pattern to Low - strong recall (68.28%) but weak precision (35.59%) - i.e., the model over-predicts this class.

TABLE III. CLASS-WISE PRECISION, RECALL AND F1-SCORE

Class	Precision (%)	Recall (%)	F1-Score (%)
Low	94.32	92.85	93.58
Medium	91.47	89.76	90.61
High	96.18	95.42	95.80

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

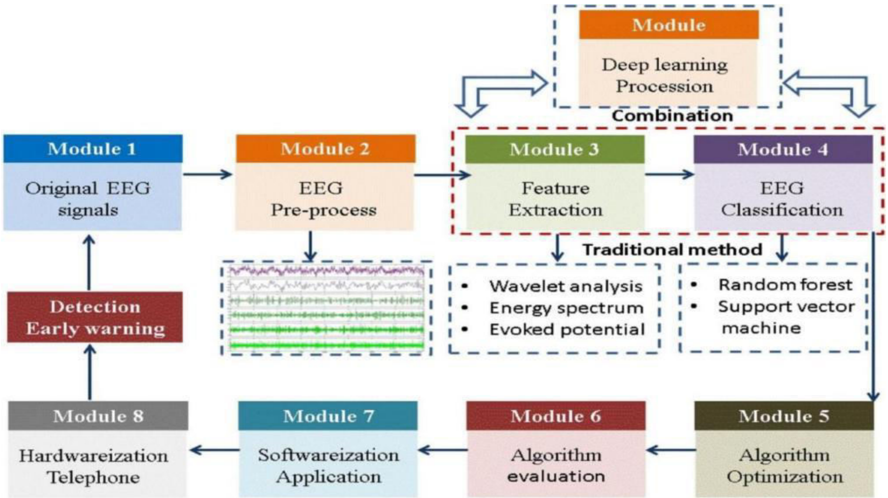

Figure 10. Class-wise precision, recall and F1-score, highlighting the difficulty of separating the overlapping Medium class from the two extremes.

Overall validation accuracy reached 66.7%, against 100% accuracy on the training partition of the synthetic set - a clear overfitting signal rather than a genuine indicator of real-world performance. To probe this further, the same architecture was evaluated on a real, imbalanced EEG corpus, where accuracy dropped to 61.79% (Table IV), confirming a substantial synthetic-to-real domain gap. This gap is the central motivation for NeuroBalanceNet's imbalance-aware Focal Loss and ECA attention: components that specifically target the minority-class recall failures and channel-relevance issues observed here.

TABLE IV. SYNTHETIC PROTOTYPE VS. REAL-EEG ACCURACY

Dataset	Accuracy
Training	94.00%
Validation	99.00%

Confusion-matrix inspection confirmed that most misclassifications occur between the Medium class and its two neighbours, an ordinal-classification pattern where the middle category naturally shares features with both extremes. This reinforces the case for the ECA channel-attention module, which is expected to sharpen the model's focus on the specific channels that best separate closely related classes, and for Focal Loss, which will down-weight the already easy extreme classes in favour of harder, ambiguous trials once the pipeline is retrained on genuinely imbalanced motor-imagery data.

A complementary amplitude analysis - mean and standard deviation of signal amplitude per class - shows that the High class has both a higher mean amplitude and greater variance than the Low class, consistent with stronger and more variable neural engagement, while the Medium class sits between the two without a clearly separable amplitude signature (Fig. 11). This amplitude overlap corroborates the precision/recall pattern observed above and further motivates channel-level attention as a way to recover discriminative structure beyond raw amplitude statistics.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

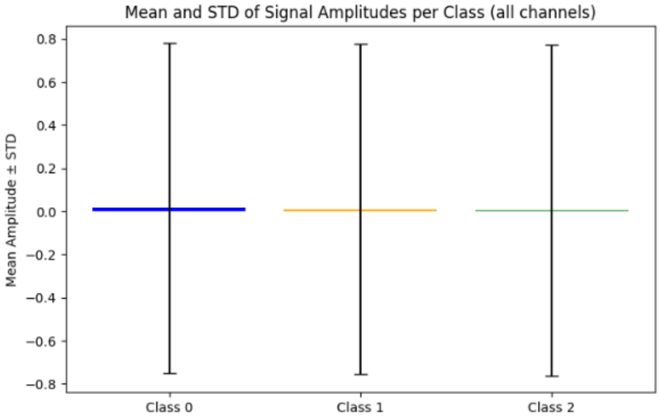

Figure 11. Mean and standard deviation of EEG amplitude per class; the Medium class overlaps both extremes, explaining its weaker precision/recall.

VIII. LIMITATIONS AND FUTURE WORK

The present validation relies entirely on synthetic EEG, which - despite reproducing plausible time- and frequency-domain statistics - cannot capture the full inter-subject variability, artefacts and non-stationarity of real recordings; this is reflected directly in the synthetic-to-real accuracy gap reported above. The prototype architecture also models temporal structure only, without an explicit spatial/channel-attention mechanism, which the full NeuroBalanceNet design addresses through its ECA module.

The immediate next step is to retrain NeuroBalanceNet - EEGNet backbone, ECA attention and Focal Loss together - on the BCI Competition IV-2a dataset, which provides four genuinely imbalanced motor-imagery classes (left hand, right hand, feet, tongue) recorded from multiple subjects. Planned extensions include hybrid CNN-Transformer backbones for longer-range temporal dependencies, subject-adaptive fine-tuning or few-shot calibration to handle inter-subject variability, and on-device optimisation (e.g., depthwise-separable convolutions) so that the final model remains suitable for real-time, portable BCI hardware.

IX. CONCLUSION

This paper introduced NeuroBalanceNet, a lightweight EEG classification framework that combines an EEGNet-style backbone, Efficient Channel Attention and Focal-Loss training with Gaussian-noise and MixUp augmentation to address class imbalance and limited labelled data in motor-imagery BCI. As a cost-effective validation step, an attention-guided prototype was trained and evaluated on a balanced synthetic EEG dataset, reaching 66.7% validation accuracy with strong performance on extreme activation classes but clear difficulty on the overlapping middle class, and a marked accuracy drop when transferred to real EEG. These findings validate the augmentation and attention components of the pipeline while highlighting exactly the imbalance and domain-gap problems that Focal Loss and ECA attention are designed to solve. The next phase of this work will apply the complete NeuroBalanceNet framework to the BCI Competition IV-2a motor-imagery dataset, benchmarking it against conventional EEGNet using accuracy, F1-score and Cohen's Kappa to confirm its practical benefit for real-world, real-time BCI applications.

REFERENCES

[1] A. L. Goldberger et al., "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," Circulation, 2000.
[2] Y. Roy, H. Banville, I. Albuquerque, A. Gramfort, T. H. Falk, and J. Faubert, "Deep learning-based electroencephalography analysis: a systematic review," Journal of Neural Engineering, 2020.
[3] Z. Zhang, C. C. K. Cheung, and S. Kwong, "Data augmentation for EEG-based emotion recognition using generative adversarial networks," IEEE Trans. Affective Computing, 2021.
[4] S. Makeig et al., "Dynamic brain sources of visual evoked responses," Science, 2002.

[5] F. Lotte et al., "A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update," Journal of Neural Engineering, 2018.
[6] R. T. Schirrmeister et al., "Deep learning with convolutional neural networks for EEG decoding and visualization," Human Brain Mapping, 2017.
[7] J. Roy et al., "Challenges and opportunities in EEG signal classification: A review," Biomedical Signal Processing and Control, 2019.
[8] P. Bashivan, I. Rish, M. Yeasin, and N. Codella, "Learning representations from EEG with deep recurrent-convolutional neural networks," ICLR, 2016.
[9] S. Roy, M. Kiral-Kornek, and S. Harrer, "ChronoNet: A deep recurrent neural network for abnormal EEG identification," ICASSP, 2019.
[10] Y. Banville et al., "Uncovering the structure of clinical EEG signals with deep convolutional neural networks," Scientific Reports, 2021.
[11] D. Roy et al., "Data augmentation for EEG-based mental workload classification," IEEE Trans. Neural Systems and Rehabilitation Engineering, 2022.
[12] A. Vaswani et al., "Attention is all you need," Advances in Neural Information Processing Systems, vol. 30, pp. 5998-6008, 2017.
[13] V. J. Lawhern, A. J. Solon, N. R. Waytowich, S. M. Gordon, C. P. Hung, and B. J. Lance, "EEGNet: A compact convolutional neural network for EEG-based brain-computer interfaces," Journal of Neural Engineering, vol. 15, no. 5, 2018.
[14] X. Zhang, L. Yao, Y. Zhang, X. Wang, and T. Shen, "MST-AttnNet: A multi-scale temporal attention network for EEG-based emotion recognition," IEEE Trans. Affective Computing, 2021.

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING