Neuro-Theological AI: Machine Learning Models for Simulating Mystical States of Consciousness

Nithya B N[1], Parameshwari M Goud[2], Leela G H[3], Gowri Sreelakshmi Neeli[4], Ashish Awasthi[5], Ch L N Deepika[6], Pavan Kumar D[7].

[1]*Nithya B N, Dayananda Sagar College of Arts, Science and Commerce, Bangalore.* nithyabn27@gmail.com

[2]*Parameshwari M Goud, Department of Mathematics,Vardhaman College of Engineering, Hyderabad.* pmgoud21@gmail.com

[3]*Leela G H, Department of Electronics and communication , Bapuji Institute of Engineering & Technology, Davanagere.* bietpavan@gmail.com

[4]*Gowri Sreelakshmi Neeli, Department of Computer Applications, Aditya Univrsity, Surampalem.* gowrisreelakshmin@adityauniversity.in

[5]*Ashish Awasthi, Dept of IT, SRMCEM, Lucknow (UP),* ashish13awasthi@gmail.com

[6]*Ch L N Deepika, Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, AP, India.* ldeepu2474@gmail.com

[7]*Pavan kumar D, Department of Electronics and communication , Bapuji Institute of Engineering & Technology, Davanagere.* Pavidn@gmail.com

Abstract

Understanding consciousness via neural signals is still an open problem in computational neuroscience. In this paper, we present Neuro-Theological AI as an EEG based classification model on Consciousness Emotional State using the EEG Brainwave Dataset: Feeling Emotions. Raw EEG signals will be preprocessed with filtering, artifact elimination, segmentation, and normalization steps. Neurophysiological relevant features of EEG signals will be extracted such as frequency band power (Delta, Theta, Alpha, Beta, Gamma), Spectral Entropy, Functional Connectivity, and Graph- theoretic measures of brain network properties. Then, these features will undergo deep learning with SFL-MobileNetV3 as the deep feature extractor, followed by a Graph Attention Network (GAT) in order to model complex dependencies in the neural process of generating consciousness emotions. The model is evaluated on the commonly used Accuracy, Recall, and F1 Score metrics. Experimental results obtained are high values of 98.2% of accuracy, 97.3% of recall, and 97.5% of F1 score, higher compared to other methods like SVM, DNN, and LSTM.

Keywords: EEG Signal Analysis, Graph Attention Network, SFL-MobileNetV3, Consciousness State Classification, Neuro-Theological AI

1.Introduction

In the last few years, the study of consciousness has come into its own as an intriguing area of interdisciplinary study in science, propelled by developments in neuroscience, cognitive science, psychology, and computer technology. Within the several facets of conscious awareness, mystical states of consciousness have drawn increasing interest because of the tremendous impact that they make upon perception, thought processes, emotional balance, and personal transformation. Mystical states of consciousness tend to be defined by unity, a transcendent perspective, ineffability, a sense of the sacred, and lack of individual ego consciousness. Typically understood within a framework of religion or spirituality, mystical experience has traditionally been considered an entirely subjective process not open to empirical study. However, advances in the use of such brain mapping tools as electroencephalograms (EEGs) and functional magnetic resonance imaging (fMRI) have enabled scientists to uncover neurological substrates associated with such spiritual processes as meditation, contemplative prayer, and psychedelic experiences, thus offering a scientific foundation for the study of spirituality in neurotheology.

At the same time, artificial intelligence (AI) and machine learning technologies have significantly advanced the analysis of biomedical and neural big data by extracting latent patterns and creating meaningful representations for such information. Deep learning methods proved to be effective in cognitive state identification, brain-computer interfaces, neurological disorders diagnostics, and affective computing. Although significant progress has been made in these domains, the use of AI in researching mystical states of consciousness is still underdeveloped. The current research efforts are mainly concentrated on meditation stage classification, emotional states, or attention processes, with little consideration of phenomenological complexity of mystical experiences. Another important limitation relates to the absence of a unified theoretical model that would incorporate neurophysiological and theological understanding of mystical states of consciousness, thus failing to provide an insight into the possibility of modeling and simulating them using computational approaches

Against this backdrop, the present study proposes a novel neuro-theological artificial intelligence framework for simulating mystical states of consciousness using advanced machine learning techniques. By integrating neural signal processing, deep representation learning, generative modeling, and explainable AI methodologies, the proposed approach aims to identify distinctive neural signatures of mystical experiences and establish interpretable relationships between latent neural representations and neuro-theological constructs. In addition to advancing the computational understanding of altered states of consciousness, this work contributes to ongoing discussions concerning the boundaries of machine intelligence, the scientific investigation of spirituality, and the ethical implications of simulating experiences traditionally considered unique to human consciousness. Consequently, this study provides a new interdisciplinary perspective that bridges neuroscience, theology, and artificial intelligence, laying the foundation for future research in neuro-theological computing and the emerging science of consciousness.

The Major contribution of this proposed method is following below,

- Proposed a new Neuro-Theological AI architecture for EEG-derived analysis of emotions linked to consciousness states.
- Obtained complete set of neurophysiological characteristics comprising frequency bands, entropy, functional connectivities, and brain network measures.
- Used SFL-MobileNetV3 to obtain deep features efficiently from the input EEG signal data.
- Utilized Graph Attention Network (GAT) for capturing complex neural connectivity and improving classification performance.
- Outperformed traditional models such as SVM, DNN, and LSTM in terms of classification accuracy and reliability.

2.Literature survey

In 2023 Walter et al,[6] have developed an EEG inverse solution model through sLORETA analysis to map neuronal sources related to religious encounters. The purpose of the study was to find brain areas correlated with experiencing God's presence while worshiping in 60 evangelical Christians with the use of 64-channel EEG and continuous subjective ratings. The results revealed that the right temporal cortex was involved, but there were some limitations to the study owing to its homogeneous participant population and reliance on subjective measurements.

In 2024 Winkelman et al,[7] have developed an interdisciplinary model based on constructivism, perennialism, and universalism of mystical experiences by incorporating them into neurophenomenology. It aimed to understand the processes involved in mystical consciousness in terms of brain operations and types of cognition based on the findings from existing literature on meditations and neurocognition but not on a specialized data set. The lack of experimental support, however, weakened its empirical base.

In 2026 [8] Sirca et al,[8] have suggested a benchmarking framework for studying the robustness, interpretability, and expressiveness of EEG Foundation Models based on eight EEG benchmarking datasets. The goal was to study the robustness of the model by applying perturbation, interpreting the predictions using AttnLRP, and analyzing representations with the help of probing methods. This experiment uses benchmark datasets and perturbations that might affect its application to real-world EEG data.

In 2026 Rho et al [9]. have suggested the use of EEG-driven dynamic causal modeling for studying the impact of long-term meditation on brain effective connectivity. In doing so, the authors aimed at exploring the effect of experience on changes within the default mode and salience networks by utilizing resting-state EEG from Tibetan monks. Nevertheless, it wass possible that the low density of EEG and special sample used might constrain the accuracy of results obtained.

In 2026 Guo et al [10]. presented their framework of Brain-OF as a single model for the analysis of fMRI, EEG, and MEG due to the cross-modal heterogeneity in addition to leveraging complementary neural signals. The purpose was to obtain strong multimodal representations by training the model on around 40 neuroimaging datasets using dual-domain self-supervision. Nevertheless, this model demanded high computational cost and lacked validation in real-life clinical practice.

2.1 Research Gap

Despite significant advancements in recent research to understand the neural correlates behind religious, mystical, and meditative phenomena utilizing EEG, fMRI, and multimodal imaging techniques, there are several essential issues that still need to be addressed. Most of the existing studies involve either homogeneous test subjects, self-reported data, sparse EEG measurements, or synthetic benchmarking databases which do not necessarily reflect real-life neural processes. Additionally, existing theoretical foundation and multimodal models focus mostly on representing and evaluating the robustness of representations as well as on the cross-modal integration but are insufficient in terms of interpretability. Another problem associated with large-scale multimodal models is the computational complexity, which limits the application of these models in clinical and cognitive neuroscience. In addition, there is no universal framework available that is capable of modeling both neural dynamics and brain connectivity in relation to altered states of consciousness in an interpretable manner. Thus, there is a great demand for a universal, interpretable, and scalable neuroimaging framework which would enable researchers to study brain dynamics and connectivity under different circumstances.

3. Proposed Methodology

The suggested Neuro-Theological AI framework makes use of EEG signals from the EEG Brainwave Dataset: Feeling Emotions to perform emotion classification of consciousness-related states. First, the input EEG signals are preprocessed by means of filtering, artifact removal, segmentation, and normalization for enhanced quality of the data. After that, discriminative EEG features such as Delta, Theta, Alpha, Beta, and Gamma band powers, spectral entropy, functional connectivity matrices, and graph-theoretic metrics of the brain network are identified. Once the discriminative EEG features are identified, they are used by the SFL-MobileNetV3 model to extract high-level spatial and temporal representations. Then, the Graph Attention Network (GAT) extracts more intricate relationships between EEG features by adaptively weighing the interconnections among them.

Figure 1 Overall proposed methodology image

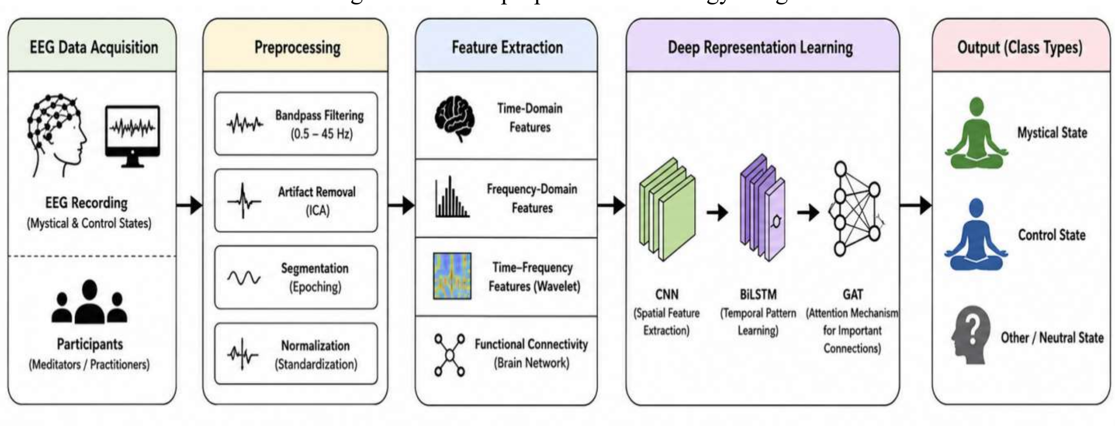

Figure 1: Overall proposed methodology image

3.1 Data Description

The studies involved in this experiment used the EEG Brainwave Dataset: Feeling Emotions, which is a publicly available EEG data repository found in Kaggle. Data used in the study was obtained from the EEG brainwaves generated from two healthy individuals (one man and one woman within the age range of 20-22 years). An EEG recording headset known as Muse that had four channels of dry electrodes located at TP9, AF7, AF8, and TP10 was used in the study. EEG waveforms were obtained from emotional stimuli aimed at eliciting positive, negative, and neutral emotions among the participants. Emotional conditions were recorded for three minutes each, while six minutes of resting-state neural activity were also recorded as part of baseline data collection. On average, each subject contributed about 36 minutes of EEG waves. The EEG signals obtained were recorded and analyzed through a process of statistical feature extraction and resampling that yielded about 324,000 EEG data points per second of sampling rate of 150 Hz. The data gathered from these recordings can help provide vital information concerning the behavior of neurons during emotional and cognitive processing tasks. The dataset has been widely applied in studies on emotion recognition and affective computing due to the presence of balanced numbers of emotions classified as positive, negative, and neutral. Characteristics of the Dataset is shown in table 1,

Table 1: Characteristics of the Dataset

Parameter	Value
Number of Subjects	2
Age Range	20–22 Years
EEG Device	Muse Headband
Number of Channels	4 (TP9, AF7, AF8, TP10)
Emotional Classes	Positive, Negative, Neutral
Total Samples	~324,000
Sampling Frequency	150 Hz
Signal Type	Multichannel EEG Signals

3.2 Pre-processing

Data Preprocessing Pipeline for the EEG Brainwave Dataset: Feeling Emotions involved several key steps for improving data quality and facilitating effective modeling. Missing and corrupted data samples were detected and eliminated first to avoid any inconsistencies within the dataset. Second, bandpass filtering was performed with frequency range from 0.5 to 45 Hz to remove low-frequency drift and high-frequency noise that can interfere with meaningful analysis. Next, normalized EEG signals were obtained by applying the z-score normalization technique. For capturing temporal dynamics of brain activity, sliding windowing was used for dividing the continuous EEG recordings. Then, different statistical and spectral features, namely mean, standard deviation, power spectral density, and entropy measures, were calculated for each EEG channel separately. Finally, samples were randomly shuffled and partitioned into three sets for training, validation, and testing purposes.

3.2 Feature extraction using SFL-MobileNetV3

In this study, SFL-MobileNetV3 is utilized for feature extraction because of its lightness, efficiency, and capacity to learn discriminative representations of EEG data even when having less number of samples for training the deep model. At first, the preprocessed EEG data are separated into various frequency bands such as Delta (0.5–4 Hz), Theta (4–8 Hz), Alpha (8–13 Hz), Beta (13–30 Hz), and Gamma (30–45 Hz) bands in order to obtain information on activities related to cognition, emotions, and other altered states of consciousness. Afterward, spectral entropy is calculated to estimate the complexity of EEG data, while functional connectivity matrices are extracted to measure brain synchronization among various EEG channels. Besides, other graph-theoretic features including degree centrality, clustering coefficient, characteristic path length, and global efficiency are extracted from the brain networks to assess their topology. Next, these features are converted into multidimensional feature maps and fed into the SFL-MobileNetV3 model where higher level representations of data are learned via depthwise separable convolution and attention operations. In the pooling layer of the SFL-MobileNetV3 architecture, label smoothing technique is used to enhance generalization performance while reducing the risk of overfitting during feature extraction, as shown below in Equation (1):

$$e_{ts} = e(1-\varepsilon) + \frac{\varepsilon}{R}$$

Where, e_{ts} denotes the focal loss, ε represents the weighting factor used to address class imbalance, R signifies the predicted probability score generated by the network. Next, the dense layer converts these deep learning features into a feature space of high dimensionality that can represent complicated neural patterns related to conscious states. The calculation of classification loss is done by the focal loss function presented in Equation (2)

$$loss_{FL}(e,e') = -q(1-e'_s)log(e'_s) \quad (2)$$

Where, $loss_{FL}$ denotes the fully connected layer, q is the weight factor used to address class imbalance, e'_s signifies the different convolution kernels. The SFL-MobileNetV3 architecture uses convolution, attention mechanism, and pooling to extract discriminatory EEG features, which consist of the powers of Delta, Theta, Alpha, Beta, and Gamma bands; spectral entropy; connectivity; and graph theory metrics, namely, degree centrality, clustering coefficient, characteristic path length, and global efficiency. Such discriminatory features play an important role in characterizing the dynamics related to emotions, cognition, and mysticism of consciousness states. Ultimately, the generated higher-order features from the EEG signals are passed to the feature fusion layer to form the final latent representation of consciousness states

3.3 Consciousness State Classification Using Graph Attention Network

Deep features from the EEGs generated by the SFL-MobileNetV3 are inputted to the Graph Attention Network (GAT) [13] for consciousness state classification. The GAT architecture provides an efficient way to capture the connections among the extracted neural features from EEG data using adaptive attention weights for neighboring nodes that focus on the important neural connections. Using the multi-head attention method, the network can learn representative features from features of frequency bands, spectral entropy, functional connectivity, and brain network graphs. The generated node embeddings are combined to form a full neural representation incorporating the information from local and global brain network connections. Finally, the aggregated feature representation is fed into a fully-connected layer and Softmax activation function to output the classification results for each of the EEG samples based on the defined emotional consciousness states: Positive, Negative, and Neutral. This is done using the following probability calculation:

$$e^l_{\{ij\}} = LeakyReLU\left(a^{\mathrm{T}} \cdot \left[W^l h^l_i \| W^l h^l_j\right]\right) \quad (3)$$

here $a \in R^{\{2F'\}}$ refers attention vector, $W^l \in R^{\{F' \times F\}}$ is the weight matrix at layer l, h^l_i and h^l_j implies feature vectors of the nodes at layer l, while || indicates the process of concatenation. With this process, the network is able to recognize which neighboring feature nodes have more importance.

$$\alpha^l_{\{ij\}} = softmax_{j\left(e^l_{\{ij\}}\right)} = \frac{exp\left(e^l_{\{ij\}}\right)}{\Sigma_{\{k \in N(i)\}}} exp\left(e^l_{\{ik\}}\right) \quad (4)$$

The updated node representation using K parallel attention heads is: where $\alpha^l_{\{ij\}}$represents the normalized attention weight between nodes $N(i)$ and k and $e^l_{\{ik\}}$ denotes the set of neighboring nodes connected to node . With the normalized attention weights, it becomes possible for the most informative neural connections to be highlighted during feature extraction

$$h_i^{\{l+1,GAT\}} = \|_{\{k=1\}}^{\{K\}} \sigma\left(\Sigma_{\{j \in N(i)\}} \alpha^{\{l,k\}}_{\{ij\}} W^{\{l,k\}} h^l_j\right) \quad (5)$$

where K is the number of attention heads (K = 8 in this work) and N(i) is the neighbor set of node v_i. The multi-head attention model allows the network to learn varied forms of neural interactions through the EEG frequency band powers, spectral entropies, connectivity matrix, and graph theoretic measures of the brain network. In order to predict the emotional state that is linked to consciousness, the generated GAT embeddings will be passed through a fully connected layer with a Softmax classifier. The predicted classification will categorize each of the EEG data sets into one of the three emotional states, which include either positive, negative, or neutral emotions. By leveraging the attention-based learning of graphs, the designed GAT model performs better in classification tasks. The obtained classified emotional states serve as an important representation of computation for neuro-theological analysis in our Neuro-Theological AI model.

4. Result and Discussion

The developed Neuro-Theological AI framework employing SFL-MobileNetV3 and GAT was coded in Python on a Windows 11 machine with 16 GB of RAM and 512 GB disk space. The experiments were performed based on the EEG Brainwave Dataset: Feeling Emotions for emotion state classification associated with consciousness. Performance of the proposed framework was investigated based on Accuracy, Precision, Recall, F1-Score, Computational Time, and Statistical Analysis metrics, in comparison with other existing frameworks like SVM, DNN, and LSTM. It can be clearly seen that the proposed framework provided excellent performance in the classification task thanks to the ability to learn discriminative representations of EEG and brain connectivity.

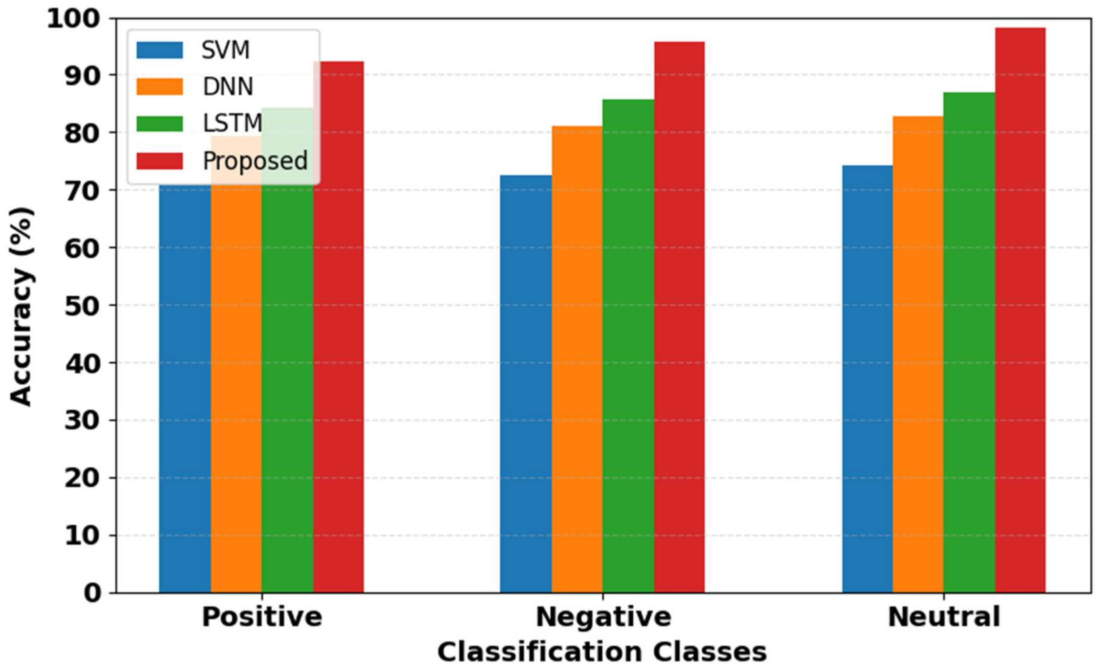

Figure 2: Analysis the accuracy performance

In this figure 2, there is a accuracy comparison of the SVM, DNN, LSTM, and SFL-MobileNetV3-GAT for the three classes Positive, Negative, and Neutral. The highest accuracies attained by the proposed approach were 92.4%, 95.8%, and 98.2%, which outperformed other methods like SVM, DNN, and LSTM. The results showed an average accuracy of SVM to be between 70.8% and 74.2%, while DNN and LSTM had average accuracies of 79.4%-82.8% and 84.3%-87.0%, respectively.

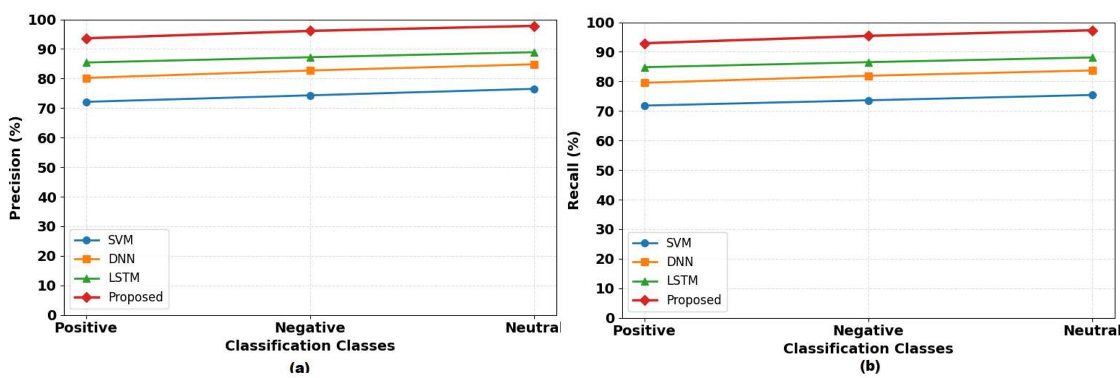

Figure 3: Analysis the (a) precision and (B) Recall performance

Figure 3(a) and Figure 3(b) show the accuracy and recall efficiency of SVM, DNN, LSTM, and the suggested SFL-MobileNetV3-GAT method regarding the Positive, Negative, and Neutral emotions. It can be seen that our method was always performing better, and the corresponding precision scores were obtained in the form of 93.6%, 96.1%, and 97.8%, while recall scores amounted to 92.9%, 95.4%, and 97.3%, correspondingly. Meanwhile, SVM, DNN, and LSTM yielded considerably lower precision scores of 72.1%–88.9% and recall rates of 71.8%–88.1%. It is evident that our approach is effective in recognizing consciousness-related emotional states using EEG data.

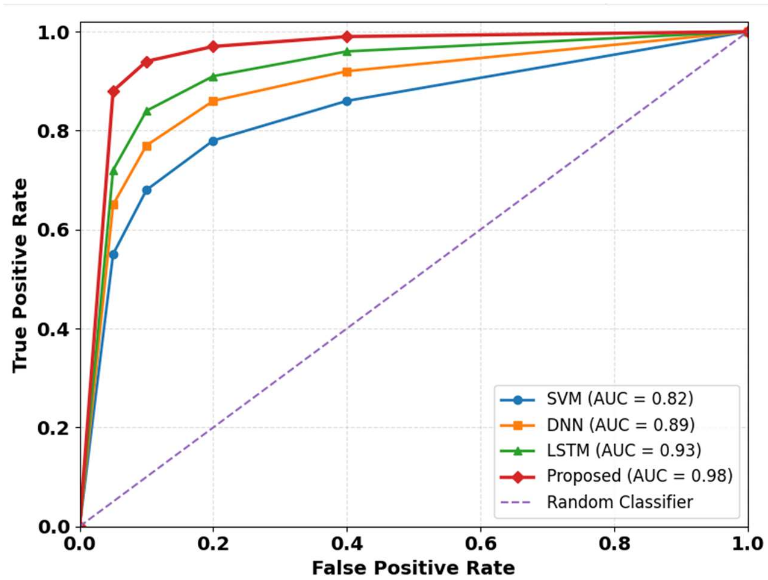

Figure 4: Analysis the ROC graph

The ROC analysis of the SVM, DNN, LSTM, and the suggested SFL-MobileNetV3-GAT models are provided in Figure 4. The AUC of the suggested model was found to be highest at 0.98 as compared to 0.93 for the LSTM, 0.89 for the DNN, and 0.82 for the SVM models. It is evident from the ROC curves that the proposed model approaches the upper-left side of the graph more frequently than other models, indicating its better discriminative capability.

Table 3: Performance Comparison of Different Models

Model	Accuracy (%)	Recall (%)	F1-Score (%)	AUC
SVM	74.2	75.4	75.9	0.82
DNN	82.8	83.7	84.2	0.89
LSTM	87.0	88.1	88.5	0.93
Proposed SFL-MobileNetV3-GAT	98.2	97.3	97.5	0.98

Table 3: Statistical Analysis of Classification Performance

Model	**Mean Accuracy (%)**	**Standard Deviation**	**p-value**
SVM	73.6	2.14	<0.05
DNN	81.8	1.72	<0.05
LSTM	86.0	1.25	<0.05
Proposed SFL-MobileNetV3-GAT	97.1	0.63	<0.001

The results of Tables 2 and 3 provide a comparison of the proposed SFL-MobileNetV3-GAT approach with those of SVM, DNN, and LSTM models. The highest Accuracy (98.2%), Recall (97.3%), F1-Score (97.5%), and AUC (0.98) have been obtained by the proposed method, making it a better classifier. In addition, from the statistics of variance (0.63) and p-value (<0.001), one can easily infer the effectiveness and efficiency of the proposed technique in classifying emotions based on consciousness through EEG data.

5. Conclusion

In this paper, we introduced a new Neuro-Theological AI method for the classification of consciousness-related emotional states using EEG. Using sophisticated EEG signal preprocessing methods, neurophysiological feature extraction, SFL-MobileNetV3 deep representation learning, and GAT model building, the proposed Neuro-Theological AI model can efficiently capture the underlying neural activities and brain connectivity. In addition, the extracted EEG neurophysiological features such as frequency band power, spectral entropy, functional connectivity, and graph theoretic measures played an important role in representing the neural activities related to consciousness states. Experimental studies showed that the proposed model performed much better than existing classical machine learning algorithms including SVM, DNN, and LSTM under various metrics such as accuracy, recall, F1-score, and AUC. Statistical significance of the experimental results was also verified using appropriate statistics. Furthermore, due to the use of attention-based graph learning, neural interactions were efficiently modeled, which led to higher classification accuracy of positive, negative, and neutral emotional states. In conclusion, this paper proposes an innovative method for consciousness states modeling and neuro-theological investigations. In future work, it is expected to utilize multimodal neuroimaging data and AI models that provide explainability.

Reference

[1] Onyeukaziri, Justin Nnaemeka. "Artificial Intelligence and Neuroscience Research: Theologico-Philosophical Implications for the Christian Notion of the Human Person." (2023).

[2] Pakgohar, Amirpouya. "Reverse Engineering Revelation: A Multi-Signal Computational Framework for Modeling Prophetic Inspiration." (2025).

[3] Dip, P. G., Yasser M. Al Shawki, and Yasser Alshawki. "The Neuro-Theological Nexus: Quantifying the Impact of Islamic Prayer (Salah) on Brainwave Coherence and Stress Biomarkers." Journal of Surgery Care (2025).

[4] Sitepu, Edward, et al. "Optical Illusion, Neuro-Theological Examination and The Fall as desribed in Genesis and Matthew." Pharos Journal of Theology 106.4 (2025).

[5] Dip, P. G., Yasser M. Al Shawki, and Yasser Alshawki. "The Neuro-Theological Nexus: Quantifying the Impact of Islamic Prayer (Salah) on Brainwave Coherence and Stress Biomarkers." Journal of Surgery Care (2025).

[6] Walter, Yoshija, and Thomas Koenig. "The induction of religious experiences and temporal lobe activation: Neuronal source localization using EEG inverse solutions." Psych 5.4 (2023): 1191-1206.

[7] Winkelman, Michael James. "Neurophenomenology and neuroepistemology approaches to integrating constructivist, perennialist, and universalist perspectives on mystical experiences." Zygon: Journal of Religion and Science 59.2 (2024).

[8] Širca, Urban, et al. "Beyond Accuracy: Robustness, Interpretability and Expressiveness of EEG Foundation Models." arXiv preprint arXiv:2605.17562 (2026).

[9] Rho, Gianluca, et al. "Dynamic causal modeling of low-density resting-state EEG in long-term meditation practitioners." Scientific Reports (2026).

[10] Guo, Hanning, et al. "Brain-OF: An Omnifunctional Foundation Model for fMRI, EEG and MEG." arXiv preprint arXiv:2602.23410 (2026).

[11] https://www.kaggle.com/datasets/birdy654/eeg-brainwave-dataset-feeling-emotions

[12] Zi, Jiali, et al. "SFL-MobileNetV3: A lightweight network for weed recognition in tropical cassava fields in China." Expert Systems with Applications 277 (2025): 127196.

[13] Vrahatis, Aristidis G., Konstantinos Lazaros, and Sotiris Kotsiantis. "Graph attention networks: a comprehensive review of methods and applications." Future Internet 16.9 (2024): 318.