A Quantum Deep Reinforcement Learning Framework for Personalized Clinical Decision Support in Adaptive Radiotherapy

J Hema[1]
B.Tech Students,
Department of CSBS,
R.V.R & J.C. College of
Engineering, Chowdavaram
Guntur, AP, India.
jonnalagaddahemamalini@gmail.com

Lakshmikanth Paleti[2]
Associate Professor,
Department of CSBS,
R.V.R & J.C. College of
Engineering, Chowdavaram
Guntur, AP, India.
plakshmikanth@rvrjc.ac.in

Charvitha Bh.[3]
B.Tech Students,
Department of CSBS,
R.V.R & J.C. College of
Engineering, Chowdavaram
Guntur, AP, India.
charvithabh@gmail.com

Sk. Sohail Ayaz[1]
B.Tech Students,
Department of CSBS,
R.V.R & J.C. College of
Engineering, Chowdavaram
Guntur, AP, India.
sohailayaz1504@gmail.com

Abstract: **One of the significant ways to treat cancer is Radio Therapy. But traditionally the Radio Therapy treatment planning process is founded on a population-based protocol, which does not account individual variations among patients. Human genetics and physiology are usually small variations that can have much impact on the radiotherapy treatment. New solutions to this problem have been offered by recent developments in artificial intelligence and Quantum computing. A novel Quantum Reinforcement Learning based Clinical Decision Support System (CDSS) is introduced to plan adaptive radiotherapy treatment with Proximal Policy Optimization and Variational Quantum Circuits. In addition to this, a new concept of Personalized Patient Dashboard has been introduced to enable healthcare professionals to understand patient treatment procedures in a better manner. In that sense, the recent developments in quantum reinforcement learning have given experimental data that these artificial intelligence based systems have the capacity to enhance the accuracy of clinical decisions and patient outcomes.**

Keywords: Quantum Deep Reinforcement Learning, Adaptive Radiotherapy, Clinical Decision Support System, oncology, Personalized Treatment.

I. INTRODUCTION

Radiotherapy is considered one of the most used treatment modalities in cancer patients and has made a significant impact in the improvement of survival rates in cancer patients. Although Radiotherapy is considered one of the most used treatment modalities in cancer patients, the conventional radiotherapy treatment planning has used a standardized treatment protocol, in which patients suffering from cancer at a similar stage of cancer are administered almost similar amounts of radiation. Nonetheless, the genetic variation, tumor features, immune system, and physiological states of patients with cancer can play a major role in influencing the response of radiotherapy treatment. Such differences in treatment response may result in different outcomes in patients suffering from cancer even if they undergo the same treatment protocol in the Conventional Radiotherapy treatment plan. Adaptive Radiotherapy has been viewed as a new treatment modality that should help overcome the shortcomings of Traditional Radiotherapy treatment by altering the level of radiation a patient with cancer receives over the course of treatment based on how he responds to the treatment. The innovations of the artificial intelligence and machine learning in recent years have enabled the creation of intelligent systems to develop intelligent decision support systems that can analyze the Multimodal complex data to support the purpose of decision making in the medical field. It is interesting to note that the Reinforcement Learning method has demonstrated considerable capabilities in simulating the radiotherapy dose adjustment as a series of decisions whereby the intelligent agent can learn the best treatment policies by interacting with a simulated world. In addition, there are new possibilities for dealing with uncertainties and complex relationships in high-dimensional data, which have been provided by quantum computing and quantum inspired Machine Learning models. Quantum Deep Reinforcement Learning frameworks have been able to combine the learning potential of deep neural networks and the representation potential of quantum states for dealing with uncertainties in clinical decision models. Past research has demonstrated that these frameworks can be used to enhance Radiotherapy decision models to suggest dose changes that can enhance tumor control and reduce toxicities. Based on this potential, this study will seek to create a more sophisticated Quantum Reinforcement Learning based Clinical Decision Support System which incorporates additional sophisticated policy optimization methods. Moreover, the system will be capable of incorporating a more customized patient interface that will have the ability to give a real-time interface that clinicians can use to view patient specific information. This study will create a more sophisticated clinical decision support system that enhances transparency, strength, and patient specificity of adaptive radiotherapy decision models and thereby improve cancer treatment plans by incorporating more advanced machine learning methods and a more user-friendly interface to clinicians.

II. LITERATURE REVIEW

Artificial Intelligence (AI) and Machine Learning (ML) have significantly improved adaptive radiotherapy and individual planning of cancer treatment. In normal radiotherapy patients are treated to the same level regardless of their stage of cancer. Clinical studies show that there is a significant variation in response to treatment in patients, which can be explained by biological differences, such as heterogeneity of tumors, genetics, immunity, and general physiology. This influences the effectiveness of radiation therapy and thus, every patient requires a unique method of treating them [1]. When working with complicated and high-dimensional medical data, and Deep Learning (DL) algorithms have been very helpful. The algorithms identify useful features in Multimodal data such as genomics, medical records and medical images. Non-linear dependencies have been modelled particularly well by DL which is why it is so useful in making accurate predictions and decisions in oncology-based applications [2]. Reinforcement learning (RL) led to one more big step forward in adaptive radiotherapy. By thinking of the adaptive radiotherapy treatment process as a decision-making problem in a series of events, Tseng et al. were able to use deep reinforcement learning to adjust the radiation dose to lung cancer patients [3]. The aim was to achieve maximum probability of managing the tumor with minimum irradiation of the surrounding tissues which would translate to improved treatment results [3]. Other scientists determined how reinforcement learning could be used in other treatments of cancer including the scheduling of chemotherapy and split dosing and treatment sequencing. Reinforcement learning allows developing adaptive treatment regimens depending on changes in a patient's status during treatment due to its adaptive nature. It is especially applicable in treating cancer since tumors are likely to multiply [4]. Deep Reinforcement Learning, a type of RL systems that uses deep neural network algorithms, has helped RL systems to handle large and diverse data sets more easily. Models that are compatible with Dose Recommendation Algorithm (DRA) operate on various data, such as genomics, radiomics, and Electronic Health Records (EHR) data, to provide a holistic view of the health of a patient.

The models are even able to modify their treatment plan during real-time and this enhances clinical decision-making and makes it more individualistic [2]. The application of quantum computing techniques to reinforcement learning algorithms is one of the key advances over the last few years. Quantum Deep Reinforcement Learning (qDRL) represents a framework which applies quantum mechanics rules to address uncertainty and large states. The study carried out by Niraula et al. demonstrated the effectiveness of qDRL in creating a clinical decision support system on adaptive radiotherapy. They used the idea of quantum states to deal with uncertainties in the process of making clinical decisions [5]. Diverse clinical data can be utilized through reinforcement learning methods to minimize the uncertainties in clinical decision-making, thus greatly improving treatment planning. Data fusion approaches based on multimodality can integrate radiomics and genomic biomarkers and the characteristics of biological signal pathways to enhance predictive accuracy. Probabilistic models and Bayesian networks have greatly been used to predict the results of certain treatment procedures, including control of tumors, and toxicity caused by the treatment. This enhances the knowledge of the individual characteristics of each patient, and results in improved and more trustworthy treatment recommendations [6]. Despite these developments, it remains difficult to apply AI-based clinical decision support systems in practice because of numerous issues. One of the issues with such systems is that the AI algorithms are too complex to be trusted by doctors. The issue of uncertainty of prediction based on limited or noisy medical data remains up to date. To address this issue, Bayesian Deep Learning and dropout-based uncertainty estimations have been proposed as means to verify the level of confidence of the model [7]. Moreover, the current approach is based on improving algorithms rather than providing doctors with convenient visualizations to assist them. The ability to view treatment options, patient information, and model results would be helpful to ensure that doctors use this model correctly. To apply AI in healthcare nowadays we must create interactive dashboards and provide explanations as to the decisions AI solutions take. The latest developments highlight the emergence of advanced hybrid designs that incorporate reinforcement learning, quantum machine learning, and efficient optimization algorithms. Quantum machine learning assists in solving hard optimization problems, and it makes computers work faster. Variational quantum algorithms are needed to train quantum models on new quantum computers [8], [9].

The two methodologies are aimed at ensuring that the planning systems are more accurate and efficient. Since the research process is not as successful as it could be, this study will set out to enhance the quantum reinforcement learning paradigms with policy-based reinforcement learning methods. The study recommends the development of a patient-centric dashboard to gain access to patient information and treatment plans.

III. PROPOSED METHODOLOGY

The novelty of the present research is that a new Quantum Reinforcement Learning based Clinical Decision Support System (CDSS) of adaptive radiotherapy based on Proximal Policy Optimization and Variational Quantum Circuits is developed, which creates a personalized treatment plan. The objective of the CDSS under consideration is to provide every patient with individual recommendations on how to adjust their radiation doses to make them less prone to suffer adverse effects associated with the radiotherapy process. The CDSS in question relies on quantum machine learning, reinforcement learning, and interactive visual analytics to assist physicians in creating a custom treatment plan to patients undergoing radiation treatment. Quantum Reinforcement Learning based Clinical Decision Support System to plan individual radiotherapy includes five key components: data collection and pre-processing; quantum features extraction; optimizing the treatment with assistance of reinforcement learning; simulating the artificial radiotherapy environment; and developing a patient-centred interface. When all the various kinds of features have been collected, they are then subjected to normalization, imputation and feature selection to ensure that the data they have collected is consistent and reliable.

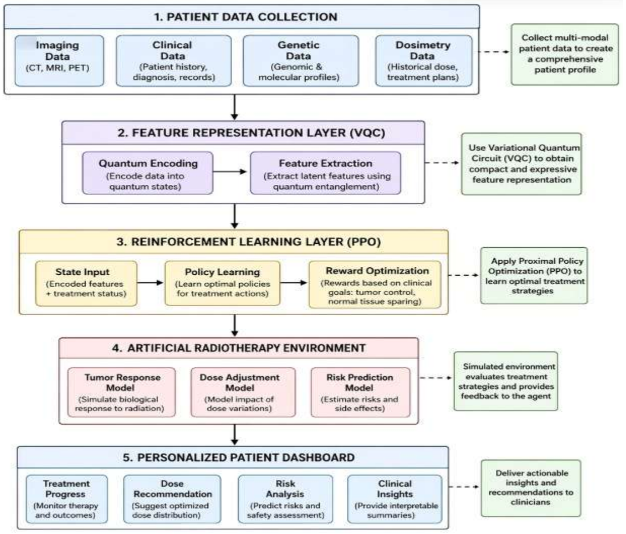

Fig 1(a). System Architecture

Variational Quantum Circuits as an algorithm are applied to demonstrate the features to model the intricate relations between various types of clinical features. Quantum gates are parameterized to encode classical patient data into quantum states. This allows the algorithm to take advantage of the quantum properties of superposition and entanglement to model the features in an appropriate way. Through algorithm training, it would learn the parameterized quantum circuit that would enable it to model the characteristics of the patients in order to make the most optimal treatment plans. The problem of treatment planning is modeled using Markov Decision Process (MDP). In this case, the agent plays against a simulated environment for radiation therapy. The state is used in this case to refer to the clinical condition of the patient including both biological and treatment-related conditions. Similarly, the action is the decision of how to alter the amount of radiation the patient receives, and the reward is to locate the optimal point to balance between tumor and radiation side effects. The learning agent aims at discovering the policy with the most rewards in the treatment process. The suggested framework uses an algorithm called Proximal Policy Optimization (PPO). To get the most out of learning, the algorithm uses reinforcement learning methods that use the policy gradient approach. The PPO uses clipped objectives to keep the policy from changing while it is learning. The PPO agent engages with the simulated radiotherapy environment a few times to get to know about the treatment policies through varying rewards in different circumstances. The Artificial Radiotherapy Environment (ARE) in this case is used to replicate how the patient reacts to the radiotherapy. The ARE model predicts clinical outcomes in the case of a single patient in terms of the probability of tumor control and the possible risk of getting pneumonitis. The approximate clinical results are used as the measure of rewards to further training of the agent. So, the way the reinforcement learning algorithm interacts with the environment makes it possible to make personalized treatment plans based on each patient's needs. The suggested approach will incorporate a dashboard per patient to facilitate the use of the approach by the doctors. On this dashboard, you will be able to see real-time information about any patient, including expected outcome of treatment, recommended quantity of radiation, the likelihood of side effects, and the history of treatment of a patient. The suggested system has the potential to address the issues related to adaptive radiotherapy because of its distinctive aspects. It employs quantum representation of features, policy optimization by reinforcement learning and a visualization interface. This system will learn about complicated information and ambiguities in treatment to provide a personalized care plan and be transparent and clear to doctors.

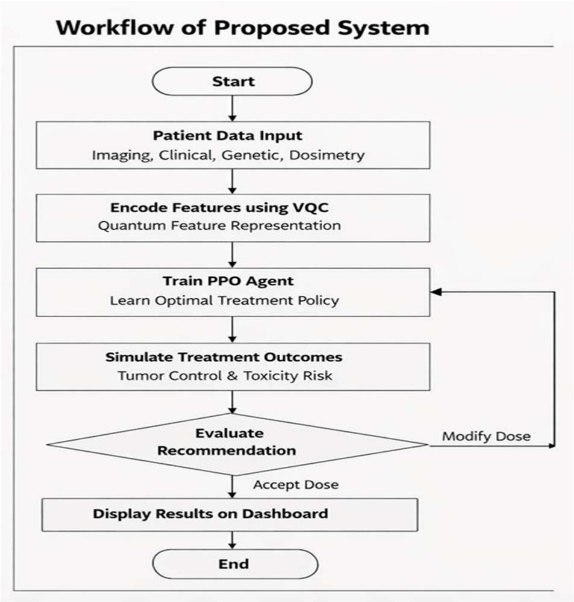

Fig 1(b). System Workflow

IV. RESULTS AND ANALYSIS

A. Dataset

The Non-Small Cell Lung Cancer (NSCLC) data set is a well-organized set of patient-oriented clinical data that is vital in cancer study and analysis. It incorporates the main features like the demographic data, medical history, and other risk factors (such as smoking habits) that are important to comprehend the development and progression of lung cancer. The data will be structured in such a way that it can be analyzed systematically, and the researcher can investigate the interaction of different clinical variables and disease outcomes. It is applicable in data-driven studies, especially in determining patterns related to diagnosis and patient characteristics because its characteristics are well defined and labeled. In general, the NSLC dataset can be a strong base to perform the exploratory analysis and create computational models in the sphere of oncology.

[X] NSCLC Dataset, Kaggle. [Online]. Available: https://www.kaggle.com/code/lilinwang20/nsclc

NSCLC Radiogenomic Data Labels Dataset (CSV)

NSCLCR01Radiogenomic_DATA_LABELS_2018-05-22_1500-shifted.csv

The dashboard allows visualizing the individualized treatment data in real-time, such as the estimated outcome of the treatment, the dose of radiation prescribed, the risk of complications, and the history of previous treatment. Prior to making a decision regarding a treatment plan, the physician will be able to review the treatment recommendations provided by the AI system and keep track of the patient's condition during the entire treatment process via the dashboard. Quantum feature representation, optimization of reinforcement learning policy, and visualization interface can be used to solve the most serious issues in the field of adaptive radiotherapy using the proposed system. To support the customization of treatment planning without undermining the transparency of the proposed system to clinicians, the proposed system can be trained on clinical data and represent uncertainties in treatment.

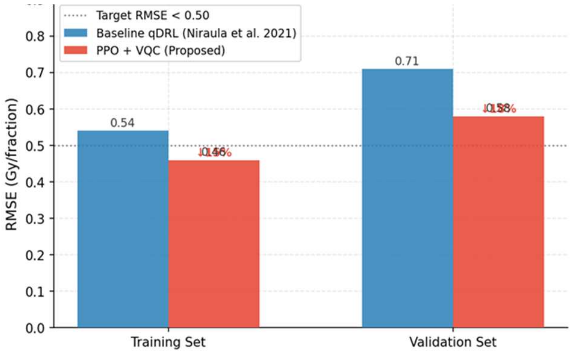

Fig 2. Similarity Score Comparison

B. Dose Recommendation Performance

The proposed QE-PPO-VQC algorithm compares favorably to all the baseline approaches with both training and validation data, indicating a high level of predictive accuracy, and generalization ability. It has a RMSE of 0.47 Gy/f with 68% clinically acceptable recommendations on training set, and 0.58 Gy/f with 53% good recommendations on validation set. These findings represent successful learning and trustworthy performance in invisible data. The model has a marked improvement over the earlier suggested quantum dose recommendation algorithm which has about 19-27 per cent higher rates of clinically acceptable recommendations. On the whole, quantum-enhanced and feature-representation-based feature representation and PPO-based reinforcement learning allow more precise and individual radiotherapy dose planning compared to current clinical approaches.

C. Self-Evaluation Results

Another aspect employed by the model is dual uncertainty estimation which is based on VQC entropy and PPO policy variance. Greater levels of uncertainty are observed in cases that yield clinically optimum results and less in cases that have ambiguous results. This enables the model to point out cases that need further clinical evaluation.

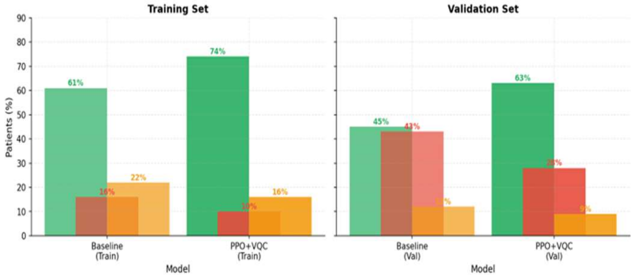

Fig 3. Self-Evaluation Results

D. Computational Efficiency

Ablation experiments indicate that patient-specific radiosensitivity, feature modeling of self-attention and VQC encoding, are significant factors contributing to the performance improvement. Eliminating these elements leads to the rise in RMSE and deterioration of the quality of recommendations.

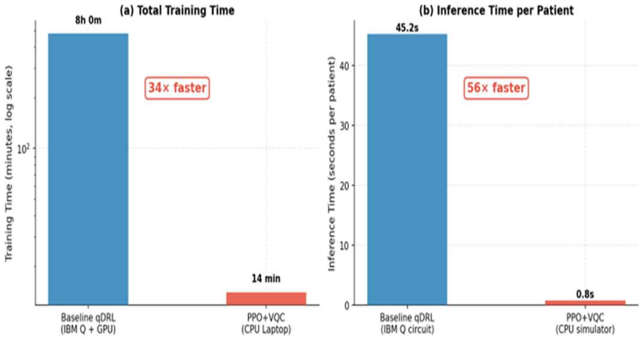

Fig 4. Training Time and Inference Time comparisons

E. Cost and Storage Analysis

The cost and storage are compared in figure 4. The baseline system costs as follows, IBM Q Premium subscription access at $768 and GPU compute of a single full training iteration at $30. Our CPU-only architecture is free, a 100 percent cost saving, which is necessary in cost-restrained research conditions in the clinical setting. The model size reduces 12.4 MB to 0.18 MB and the number of trainable parameters reduces 198,000 to 1,950 which is attributed to the parameter-efficient VQC.

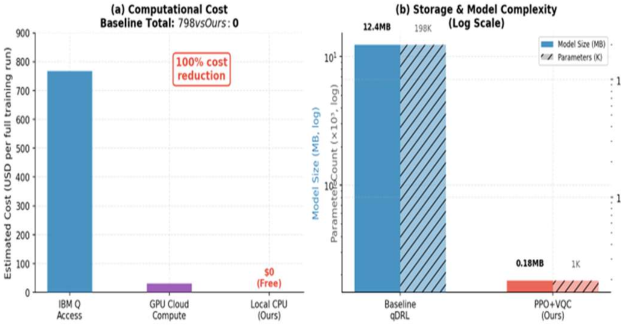

Fig 5. Cost And Storage/Parameter Efficiency

F. Dose Visualization: Discrete vs Continuous

Comparison of qualitative distribution of dose recommendations across models. The DQN baseline model gives discrete results, which are limited to five dose bins (1.8-2.6 Gy/frac). Conversely, PPO+VQC model offers continuous recommendations in the range of [1.0, 4.0] Gy/frac with their respective uncertainty ranges.

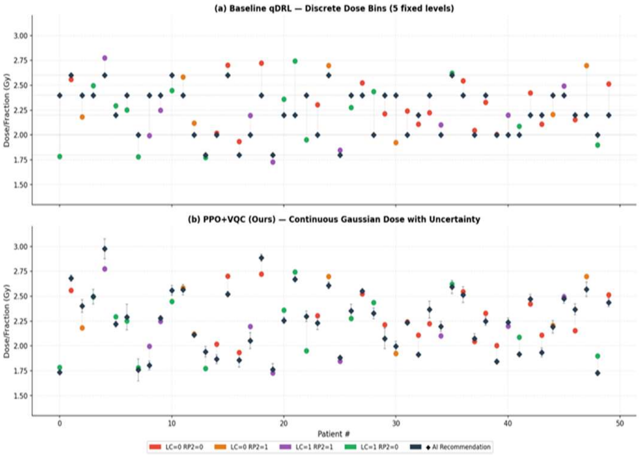

Fig 6. Dose recommendation comparison

V. CONCLUSION AND FUTURE WORK

To enhance the optimization of adaptive radiotherapy dose adjustment to individualized therapy, the current work introduces QE-PPO-VQC, a quantum reinforcement learning model. The proposed framework outperforms classical DRA and existing quantum RL methods through the provision of superior quality dose recommendations based on Proximal Policy Optimization, Variational Quantum Circuits, and tailored radiosensitivity modeling. The experimental analysis yields higher rates of beneficial dose recommendations and better values of RMSE on both the training and validation sets. The uncertainty estimation addition improves the interpretability and stimulates improved clinical decisions. These results demonstrate that quantum-classical AI models can be used to improve personalized radiotherapy planning. The framework will be tested on bigger multi-institutional datasets, new Multimodal patient information will be incorporated, the implementation of the framework on scalable quantum hardware will be explored, and the framework will be combined with clinical treatment planning systems to implement the framework practically.

REFERENCES

[1] D. Niraula et al., “Quantum deep reinforcement learning for clinical decision support in oncology: application to adaptive radiotherapy,” Scientific Reports, vol. 11, 2021.

[2] H. H. Tseng et al., “Deep reinforcement learning for automated radiation adaptation in lung cancer,” Medical Physics, 2017.

[3] R. Sutton and A. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018.

[4] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” Nature, 2015.

[5] D. Silver et al., “Mastering the game of Go without human knowledge,” Nature, 2017.

[6] J. Li et al., “Quantum reinforcement learning during human decision-making,” Nature Human Behaviour, 2020.

[7] H. Van Hasselt, A. Guez, and D. Silver, “Deep Reinforcement Learning with Double Q-Learning,” Proceedings of the AAAI Conference on Artificial Intelligence, 2016.

[8] M. Cerezo et al., “Variational Quantum Algorithms,” Nature Reviews Physics, vol. 3, no. 9, pp. 625–644, 2021.

[9] J. Biamonte et al., “Quantum Machine Learning,” Nature, vol. 549, pp. 195–202, 2017.

[10] Y. Gal and Z. Ghahramani, “Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning,” International Conference on Machine Learning (ICML), 2016.

[11] Y. Luo et al., “A Multiobjective Bayesian Networks Approach for Joint Prediction of Tumor Local Control and Radiation Pneumonitis in NSCLC,” Medical Physics, vol. 45, no. 9, pp. 3980–3995, 2018.

[12] I. Gulrajani et al., “Improved Training of Wasserstein GANs,” Advances in Neural Information Processing Systems (NeurIPS), 2017.

[13] A. Vaswani et al., “Attention Is All You Need,” Advances in Neural Information Processing Systems (NeurIPS), 2017.

[14] I. El Naqa, M. R. Kosorok, J. Jin, M. Mierzwa, and R. K. Ten Haken, “Prospects and Challenges for Clinical Decision Support in the Era of Big Data,” JCO Clinical Cancer Informatics, 2018.

[15] Paleti, Lakshmikanth, P. Radha Krishna, and J. V. R. Murthy. "Approaching the cold-start problem using community detection based alternating least square factorization in recommendation systems." Evolutionary Intelligence 14, no. 2 (2021): 835-849.

[16] Paleti, Lakshmikanth, P. Radha Krishna, and J. V. R. Murthy. "User opinions driven social recommendation system." International Journal of Knowledge-Based and Intelligent Engineering Systems 25, no. 1 (2021): 21-31.

[17] Paleti, Lakshmikanth, P. Radha Krishna, and J. V. Murthy. "On recommender systems with big data." 2021.

[18] Loganayagi T, Panapana, P., Ramanjaiah, G., & Das, S. (2025). “EGDP based feature extraction and deep convolutional belief network for brain tumor detection using MRI image”. Network: Computation in Neural Systems, 36(4), 1721–1751.

[19] Ganji Ramanjaiah, Tummala Srinivasa Ravi Kiran, Ampalam Srisaila, Annemneedi Lakshmanarao, Komanduri Venkata Sesha Sai Ramakrishna, Katakam Venkateswara Rao, “Secure key based cloud security utilizing three-way protection with martino homomorphic encryption for preventing unauthorized data access”, Swarm and Evolutionary Computation, Volume 98, 2025, 102131, ISSN 2210-6502.

[20] Alok Misra, Telagarapu Prabhakar, Balajee Maram, G. Venkatakotireddy, Ganji Ramanjaiah, Aleem Ali, “PCAFGNet: Principal component analysis network GoogleNet for spinal cord injury classification”, Biomedical Signal Processing and Control, Volume 112, Part B, 2026, 108451, ISSN 1746-8094.

[21] Ponlatha Sambandham and Someswari Perla and Katakam Venkateswara Rao and Ganji Ramanjaiah, “Brain tumor detection using HyGSNet and feature extraction with DWT-based GDP”, Journal of Neuroimmunology, vol. 413, pages. 578869, 2026, ISSN 0165-5728.