

Hybrid Deep Learning and Reinforcement Learning Approach for Brain Tumor Classification from MRI Images

Çiğdem Gökçek-Saraç^{1,*}, Seda Arıkan Arıbal², Yiğit Ali Üncü³

¹Department of Biomedical Engineering, Faculty of Engineering, Akdeniz University, 07050, Antalya, TÜRKİYE

<https://orcid.org/0000-0002-3538-6551>

*corresponding author: ciqdemsarac@akdeniz.edu.tr

²Department of Biomedical Engineering, Faculty of Engineering, Akdeniz University, 07050, Antalya, TÜRKİYE

<https://orcid.org/0009-0009-9067-0502>

³Vocational School of Technical Sciences, Department of Biomedical Equipment Technology, Akdeniz University, 07050, Antalya, TÜRKİYE

<https://orcid.org/0000-0001-7398-9540>

(Received: 07.05.2025, Accepted: 14.11.2025, Published: 26.11.2025)

Abstract

Brain cancer, resulting from abnormal tumor growth in brain tissue, requires accurate and timely diagnosis. Although MRI plays a crucial role, manual interpretation is prone to errors and delays. To address this, we propose a hybrid system combining deep learning (VGG16, ResNet50, DenseNet201) with reinforcement learning (Q-learning) for brain tumor classification. Using three distinct MRI datasets within MATLAB, the models achieved high classification accuracies: 97.45% (VGG16), 96.06% (ResNet50), and 96.93% (DenseNet201). The integration of reinforcement learning improved decision-making and interpretability. Additionally, a user-friendly interface was developed to support clinical decision-making. This study demonstrates that combining deep learning with reinforcement learning enhances model adaptability, offering a more reliable and effective diagnostic approach.

Keywords: Brain tumor, MRI classification, Deep learning, Reinforcement learning, CNN models

1. Introduction

The brain, as the central nervous system's control center, governs motor control, cognition, and vital physiological processes. Cancer, a result of abnormal cell growth, can affect any tissue, but brain tumors are among the most severe and life-threatening forms [1]. According to the CBTRUS report (2015–2019), brain and CNS tumors caused 84,264 deaths in the U.S., with 93,470 new cases reported of which 26,670 were malignant [2].

Common adult brain tumors include gliomas, meningiomas, and pituitary tumors [3]. Gliomas, particularly glioblastomas, are highly malignant, while meningiomas are typically benign and slow-growing. Pituitary tumors, though usually benign, can disrupt hormonal balance [4]. Due to their biological complexity, traditional treatment methods often fall short. Therefore, improvements in early diagnosis, biomarker identification, and AI-supported medical imaging are critical. Though neurological assessments are useful, advanced imaging techniques such as CT and MRI are essential for definitive diagnosis [5].

MRI is preferred for its superior soft tissue contrast, spatial resolution, and absence of ionizing radiation. Gadolinium-based contrast agents further enhance visualization of tumor vascularity [6,7]. However, manual interpretation of MRIs remains time-consuming and susceptible to observer variability [8]. AI and ML techniques promise to improve diagnostic accuracy and reduce human error [9].

AI, emerging post-1950, encompasses machine learning and problem-solving. Initially limited by hardware constraints, its progress accelerated in the 1990s [10]. Deep learning, particularly CNNs, has advanced medical imaging through classification, segmentation, and feature extraction [11]. Architectures such as VGG, ResNet, and DenseNet are widely adopted [12,13]. Reinforcement learning (RL), through trial-and-error learning, complements CNNs by dynamically adapting classification strategies using Q-learning [14].

Combining RL with CNN architectures such as VGG16, ResNet50, and DenseNet201 represents a significant advancement over static CNN models, enabling continuous adaptation and improved generalization [15,16]. This integration mitigates overfitting and enhances performance on diverse datasets. Despite their promise, CNNs often struggle with generalization and overfitting when applied to limited or imbalanced datasets [17]. These limitations are particularly pronounced in brain tumor imaging. This study addresses these challenges by proposing a reinforcement learning (RL)-augmented convolutional neural network (CNN) model that dynamically learns from environmental feedback, leading to improved classification performance. Recent research highlights the potential of hybrid frameworks combining deep learning and RL for robust clinical decision-making [18].

Interpretability remains a significant challenge. CNN models, often criticized as black boxes, hinder clinical trust and widespread application. Although supervised learning achieves high accuracy, it is heavily dependent on large labeled datasets, which is a notable drawback [19]. In contrast, reinforcement learning enhances model adaptability through sequential learning, although its application to brain tumor diagnosis remains limited [20]. Existing literature still predominantly focuses on supervised deep learning methods, with few studies integrating RL with CNNs like VGG16, ResNet50, or DenseNet201 for brain tumor classification. However, promising results using hybrid models and CNN architectures in brain tumor imaging have started to emerge [21-23].

To address this gap, this study proposes a hybrid AI framework for classifying primary brain tumors using CNNs integrated with Q-learning. Current diagnostic approaches often depend on subjective assessments by radiologists. While CNNs offer automation, their limited generalizability restricts clinical deployment. The proposed model leverages the strengths of VGG16, ResNet50, and DenseNet201 within a reinforcement learning framework to overcome issues associated with small and imbalanced datasets. This approach aims to enhance diagnostic consistency, reduce clinician variability, and support personalized treatment planning, ultimately improving clinical decision-making and patient outcomes. This study makes several key contributions to the field of medical image classification. It introduces a hybrid deep-learning framework that integrates reinforcement learning (Q-learning) with well-known CNN architectures (VGG16, ResNet50, and DenseNet201) for brain tumor classification. The proposed approach formulates an interpretable reinforcement structure by explicitly defining states, actions, and reward functions linked to classification performance. Through this design, the model achieves high classification accuracy (up to 97.95%) and robust performance across multiple datasets while reducing overfitting. Furthermore, the framework provides an adaptable and reproducible implementation that can serve as a foundation for future clinical decision-support systems.

2. Material and Method

2.1. Dataset and preprocessing

This study utilized three publicly available brain tumor MRI datasets [24-26], each including four categories: glioma, meningioma, pituitary tumor, and healthy tissue. A total of 7168 MRI images were collected, annotated, and prepared for classification using deep learning models. All images were standardized by converting them to RGB format and resized to 224×224 pixels using MATLAB's `augmentedImageDatastore` function to ensure compatibility with the selected CNN architectures. Pixel values were normalized to the [0,1] range using the `im2double` function to facilitate efficient training. These preprocessing steps aimed to provide uniformity in data input and to optimize model learning.

2.2. Data splitting

The preprocessed dataset was randomly divided into training (70%), validation (15%), and testing (15%) subsets using MATLAB's `splitEachLabel` function, ensuring balanced class distribution across all sets.

- The training set was used to optimize model parameters.
- The validation set supported hyperparameter tuning and overfitting prevention.
- The test set, unseen during training, provided an unbiased evaluation of model performance.

2.3. Data augmentation

To improve model generalization and mitigate overfitting, data augmentation was applied to the training set. Techniques included random rotations (0°-30°), brightness adjustments (80%-120%), horizontal flipping, and random cropping. These operations simulated real-world variability and enhanced the model's robustness to spatial and illumination differences across MRI scans.

2.4. Deep learning model development

Three pre-trained CNN architectures such as VGG16, ResNet50, and DenseNet201 were employed through transfer learning within MATLAB's Deep Learning Toolbox. Each model was initialized with ImageNet weights and fine-tuned for four-class classification by modifying the final fully connected and softmax layers.

- VGG16 uses 3×3 convolution filters in a deep, uniform architecture, ideal for detailed feature extraction.
- ResNet50, with 50 layers and residual connections, addresses vanishing gradient issues and supports deeper training.
- DenseNet201 connects each layer to all preceding layers, promoting feature reuse and reducing overfitting with fewer parameters.

All models were trained on input images resized to 224×224 pixels, aligning with their architectural requirements.

2.5. Supervised learning approach

Labeled MRI scans were used to train each CNN model under a supervised learning framework. During training, the models learned to associate input features with corresponding class labels using the Stochastic Gradient Descent with Momentum (SGDM) optimizer. The validation set was employed to monitor training progress and

adjust hyperparameters, while the test set was used for final performance evaluation. Metrics such as accuracy, precision, recall, and F1-score were calculated to assess classification success and generalization ability.

2.6. Reinforcement learning and Q-table approach

To improve the classification performance of the deep learning models, a reinforcement learning-based Q-table approach was applied. This method allowed the estimation of class probabilities and provided deeper insights into prediction uncertainties, enhancing the interpretability of model outputs. The Q-learning algorithm was configured with 16 states, 4 actions, a learning rate (α) of 0.1, a discount factor (γ) of 0.9, and an epsilon-greedy strategy (initial epsilon = 1.0, minimum epsilon = 0.01, decay rate = 0.99). Training was conducted over 500 episodes with 10 steps per episode to ensure a stable balance between exploration and exploitation (Table 1).

Table 1. Hyperparameter settings of the Q-learning algorithm utilized in the proposed brain tumor classification model

Number of States	Number of Actions	Learning Rate (α)	Future Reward Discount Factor (γ)	Initial Epsilon	Minimum Epsilon	Epsilon Decay Rate	Number of Episodes	Number of Steps per Episode
16	4	0.1	0.9	1.0	0.01	0.99	500	10

During training, the Q-table was updated based on rewards received for each state-action pair. After convergence, it was adapted to predict class probability distributions. The final Q-tables for VGG16, ResNet50, and DenseNet201 models were compared using heatmaps, and an interactive interface was developed to visualize model performance and support model selection.

In this study, the state space was defined as 16 distinct conditions representing the distribution of class probabilities derived from CNN outputs. The action space consisted of 4 actions corresponding to the four target classes (glioma, meningioma, pituitary, and no tumor). The reward function was designed to assign a positive reward (+1) for correct classification and a penalty (−1) for incorrect classification. This ensured that the Q-learning agent gradually optimized its policy toward accurate tumor identification. The Q-table was updated using the Bellman equation with $\alpha = 0.1$ and $\gamma = 0.9$, balancing short-term and long-term rewards. This explicit formulation enhances the reproducibility of our approach.

While the Q-tables provide a clear mapping of learned policies and decision transitions, they should not be interpreted as visual explanations of CNN behavior. Instead, they reveal how the reinforcement agent updates its action policy based on state-reward interactions. Accordingly, the interpretability of our system lies in understanding decision logic rather than feature localization. We have revised our claims to reflect this distinction, clarifying that Q-tables primarily support transparency of decision policies, whereas true visual interpretability would require feature-level explanation methods such as Grad-CAM or LIME. The methodological novelty of this work lies in the adaptive integration of Q-learning with CNN architectures, where the reward function is explicitly tied to classification accuracy and penalizes incorrect predictions. This design enables dynamic decision-boundary adjustment and greater learning stability, which collectively enhance the model's adaptability to heterogeneous MRI data.

To clarify the connection between CNN feature outputs and Q-learning updates, an additional explanatory description was incorporated. The CNN generates probability distributions that serve as the state representations for the Q-learning agent. These states are used to update the Q-table through reward feedback, linking the CNN's probabilistic outputs to reinforcement-based decision refinement. This textual clarification was provided in place of a schematic figure, as the workflow is conceptually straightforward.

3. Results

3.1. Dataset preparation

The dataset consisted of 7168 brain MRI scans categorized into glioma, meningioma, pituitary tumor, and healthy (no tumor) classes (Table 2). All images were resized to 224×224 pixels and converted to RGB format in MATLAB to match the input requirements of the VGG16, ResNet50, and DenseNet201 architectures.

Table 2. Distribution of MRI scans by tumor type

Tumor Type	Number of Images
Glioma	1548
Meningioma	1362
Pituitary	1787
No Tumor	2471

The dataset was partitioned into training (70%), validation (15%), and testing (15%) sets using MATLAB's `splitEachLabel` function, maintaining balanced representation across classes (Table 3).

Table 3. Dataset splitting across tumor types

Tumor Type	Training	Validation	Testing
Glioma	1084	232	232
Meningioma	954	204	204
Pituitary	1251	268	268
No Tumor	1729	371	371

Data augmentation techniques, including random rotation (0°-30°), brightness adjustments (80%-120%), and horizontal flipping, were applied to the training set to improve model generalization. These augmentations exposed models to various spatial and illumination variations, thereby enhancing classification robustness.




Figure 1. Preprocessing of brain MRI images: (A) Original grayscale T2-weighted MRI; (B) Pseudo-colored version to enhance contrast and feature distinction for improved tumor detection and classification by CNNs [27].

Brain MRI images are preprocessed and transformed to serve as input for deep learning algorithms. The original grayscale T2-weighted MRI image (A) is converted into a pseudo-colored representation (B) to enhance contrast and feature distinction during training. This color enhancement improves feature extraction by convolutional neural networks, leading to more accurate tumor detection and classification (Figure 1).

3.2. Model training and performance evaluation

3.2.1. VGG16 model

The VGG16 model was trained using the same hyperparameters as summarized in Table 4. It achieved the highest validation accuracy (97.95%) and F1-score (0.9759), with the longest training time. The training and validation curves are shown in Figure 2.




Figure 2. Training effectiveness of the proposed reinforcement learning-integrated VGG16

3.2.2. ResNet50 model

The ResNet50 model also underwent 10 epochs of training with similar hyperparameters. Validation accuracy reached 96.09%, and the F1-score was 0.95437. The model exhibited rapid convergence with low loss values (Figure 3), maintaining stability across training and validation datasets. Training time was significantly shorter at 203 minutes, though slight misclassifications, particularly within the meningioma class, were observed.




Figure 3. Training performance of a ResNet50-based model combined with reinforcement

3.2.3. DenseNet201 model

The model demonstrated smooth convergence (Figure 4) and consistent performance across training and validation datasets. Training took 425 minutes, indicating a balance between classification performance and computational efficiency. DenseNet201 exhibited strong performance with fewer parameters, supporting its potential for resource-limited environments.




Figure 4. Training results for a DenseNet201-based model enhanced with reinforcement

3.2.4. Comparative performance analysis

All CNN architectures (VGG16, ResNet50, and DenseNet201) were trained under identical hyperparameter settings using transfer learning. The key training parameters and performance results are summarized in Table 4. This tabular format minimizes repetition and allows direct comparison of model efficiency and accuracy. While VGG16 achieved the highest validation accuracy, ResNet50 offered the shortest training time. DenseNet201 provided a favorable balance between the two.

An analysis of variance (ANOVA) was performed on the F1-scores of the three models, yielding an F-statistic of 0.661 and a p-value of 0.540 ($p > 0.05$), confirming that there were no statistically significant differences between the models.

Table 4. Summary of training settings and performance of CNN architectures

Model	Epoch	Optimizer	Learning Rate	Mini-Batch Size	Validation Accuracy	F1-Score	Training Time (min)
VGG16	10	SGDM	0.0001	32	97.95%	0.9759	561
ResNet50	10	SGDM	0.0001	32	96.09%	0.9544	203
DenseNet201	10	SGDM	0.0001	32	96.93%	0.9634	425

3.3. Confusion matrix analysis

Confusion matrices for all three CNN architectures demonstrated high classification accuracy across classes, with minor misclassifications in the meningioma class. Figures 5–7 present the confusion matrices for VGG16, DenseNet201, and ResNet50, respectively. While VGG16 achieved near-perfect results, DenseNet201 performed consistently across all classes, and ResNet50 exhibited slightly higher misclassification in meningioma




Figure 5. Confusion matrices illustrating the performance of the reinforcement learning-supported deep learning model: (A) Test dataset results show near-perfect classification for glioma, pituitary, and no tumor images, with minor errors in the meningioma class; (B) Training dataset results demonstrate high accuracy across all classes, indicating strong generalizability and reliability in multi-class brain tumor detection.




Figure 6. Confusion matrices showing the classification performance of the DenseNet201 model with reinforcement learning: (A) Test results demonstrate high accuracy, especially for glioma, pituitary, and no tumor classes, with minor errors in meningioma; (B) Training results confirm strong overall performance with minimal class confusion.




Figure 7. Confusion matrices of the ResNet50 model with reinforcement learning: (A) Test results show high accuracy in identifying glioma, pituitary, and no tumor cases, with slight misclassification in meningioma; (B) Training results confirm consistent performance across all classes, supporting the model's reliability in distinguishing brain tumor types and healthy cases.

Overall, all three CNN models, enhanced with reinforcement learning, demonstrated strong classification capabilities for multi-class brain tumor detection. VGG16 achieved the highest accuracy and F1-score but required the longest training time. DenseNet201 offered a favorable trade-off between accuracy and computational cost, performing consistently across all tumor types. ResNet50 achieved fast training and competitive results but showed slightly lower precision in distinguishing meningioma instances. These results highlight the importance of balancing model accuracy, generalization ability, and computational efficiency when designing AI-based brain tumor classification systems.

Across all models, minor misclassifications were observed predominantly in the meningioma class. This tendency can be attributed to the high intra-class variability of meningioma MRI appearances and the overlap of their visual features with other tumor types, as reported in previous studies. These intrinsic challenges explain the relatively lower classification precision for meningioma compared to glioma and pituitary tumors.

3.4. Q-table construction and prediction

In this study, a reinforcement learning-based Q-table approach was employed to enhance the classification process and present model predictions as percentage

probabilities. This method enabled more interpretable, probability-based outputs across the four brain tumor classes, improving the transparency of the classification results (Figures 8-11).

The Q-table was designed to optimize decision-making by representing 16 distinct classification states and four corresponding actions (tumor classes). The reinforcement learning algorithm was configured with a learning rate (α) of 0.1 to ensure gradual model adaptation and a future reward discount factor (γ) of 0.9 to emphasize long-term decision benefits.

An epsilon-greedy strategy was adopted to balance exploration and exploitation, with the initial epsilon set at 1.0, decreasing by 1% at each step (decay rate: 0.99) until reaching a minimum threshold of 0.01. Training was conducted over 500 episodes, with a maximum of 10 steps per episode, enabling the Q-table to converge toward optimal classification policies. These hyperparameter settings were crucial in refining the Q-table, enhancing both the stability and reliability of model predictions. In particular, the use of the epsilon-greedy approach was instrumental in maintaining an effective balance between the exploration of new actions and the exploitation of known optimal policies. As a result, the classification outcomes became not only accurate but also more interpretable and trustworthy from a clinical perspective.

The developed brain tumor detection system features a graphical user interface (GUI) that integrates deep learning and reinforcement learning methodologies. Users can upload MRI images and receive classification results from VGG16, ResNet50, and DenseNet201 models, including predicted tumor type, classification accuracy, and confidence scores (Figure 8-11).




Figure 8. The interface output for a glioma MRI shows accurate classification by all models with high confidence scores (VGG16: 100%, ResNet50: 99.8983%, DenseNet201: 99.9865%), supported by Q-tables illustrating reinforcement learning-based decision mapping.




Figure 9. The interface output for a pituitary tumor MRI shows accurate classification by all models with high confidence scores (VGG16: 99.9999%, ResNet50: 99.5939%, DenseNet201: 99.9669%), supported by Q-tables illustrating reinforcement learning-based decision mapping.




Figure 10. The system correctly identifies a healthy (non-tumor) brain MRI across all models with high confidence (VGG16: 99.9996%, ResNet50: 99.5772%, DenseNet201: 99.8954%). The corresponding Q-tables visualize the learned decision strategies.




Figure 11. For a meningioma case, VGG16, ResNet50, and DenseNet201 models accurately classify the MRI with high confidence, with Q-tables demonstrating the reinforcement learning-enhanced interpretability and decision pathways.

To comprehensively assess the performance of the proposed hybrid model, its results were compared with several recent studies that employed deep learning or hybrid frameworks for brain tumor classification. Table 5 presents a summary of these comparisons, highlighting the datasets, model architectures, classification categories, and reported accuracies. This comparative evaluation aims to position the proposed CNN + RL model within the current state-of-the-art and to emphasize its advantages in terms of adaptability and generalization.

Table 5. Comparative analysis of the proposed hybrid CNN + RL model and recent studies in the literature on brain tumor classification

Study	Model / Method	Dataset	Classes	Accuracy (%)	Key Characteristics
This study	VGG16+Q-Learning	ThreeMRI datasets(7168 images)	4(Glioma, Meningioma, Pituitary, Healthy)	97.95	Hybrid CNN + RL framework, dynamic decision boundary adjustment, MATLAB GUI
Pilaoon et al. (2024) [32]	Transfer Learning (VGG16, ResNet50)	Glioma Dataset	3	99.19	Transfer learning without RL, high accuracy but limited class scope
Khan et al. (2025) [33]	AlexNet, MobileNetV2, GoogleNet	BrainMRI Dataset	3	96–98	Traditional CNN transfer learning, no reinforcement layer
Munira& Islam(2022) [30]	HybridDeep Learning	Brain MRI	4	95.67	Ensemble deep model, no dynamic learning adaptation
Neamah etal. (2024) [29]	Improved ResNet50	BraTS	3	96.80	Fine-tuned ResNet50, limited interpretability
Amou etal. (2022) [23]	CNN + Bayesian Optimization	Brain MRI	4	94.20	Bayesian hyperparameter tuning
Sevinc et al. (2025) [34]	Transformer+ Distillation	Medical Images	Multi-class	98.00	Limited-data optimization via model distillation

4. Conclusion

In this study, we developed a hybrid brain tumor classification system that integrates deep learning (CNN models: VGG16, ResNet50, DenseNet201) with reinforcement learning (Q-learning) techniques, aiming to enhance diagnostic accuracy and generalizability across MRI datasets. Our experimental results demonstrated that the hybrid models achieved high classification accuracies with VGG16 yielding the highest performance (97.95%), followed closely by DenseNet201 (96.93%) and ResNet50 (96.09%). Compared to traditional supervised deep learning models, which rely heavily on static datasets and suffer from issues like overfitting and lack of adaptability [28], our reinforcement learning-augmented framework demonstrated superior adaptability and robustness. Particularly, Q-learning enhanced the models' ability to adjust decision boundaries dynamically, thus mitigating the risk of overfitting even when trained on limited or imbalanced datasets a common challenge in medical imaging studies [29].

Our results are consistent with prior studies advocating for reinforcement learning integration in medical imaging [30], but extend the literature by focusing specifically on brain tumor classification, an area where the application of RL remains relatively underexplored [31]. By coupling Q-learning with CNN architectures, we achieved both high predictive performance and increased model resilience against variations in data distribution, an important criterion for clinical applicability.

Our results (96–98 % accuracy) are comparable with or higher than those reported in recent literature. For instance, Pilagoon et al. [32] achieved up to 99.19 % accuracy using transfer learning in glioma classification. The study by Khan et al. [33] also reports high accuracy across multiple CNN architectures (AlexNet, MobileNetV2, GoogleNet) for brain tumor classification. These findings indicate that our hybrid CNN + RL approach provides competitive performance within the current state of the art.

Recent works from 2024-2025 also provide important context for our study. Sevinc, Uçan and Kaya [34] introduced a distillation approach to transformer-based medical image classification with limited data, and showed that applying distillation to transformer models yielded accuracy improvements of 1-2 %. Although their focus was transformers, the principle of enhancing a backbone model with an additional learning layer is conceptually similar to our reinforcement learning integration.

Taken together, these recent studies strengthen the rationale behind layered or hybrid learning frameworks in medical imaging and suggest that our CNN + RL integration is aligned with current trends toward enhancing performance and adaptability.

Moreover, among the three CNN backbones evaluated, VGG16 exhibited the highest classification performance, albeit with longer training times. DenseNet201, on the other hand, offered a favorable balance between computational efficiency and classification accuracy, making it a promising candidate for deployment in resource-constrained clinical settings.

In conclusion, this study provides strong evidence that the integration of reinforcement learning with deep learning architectures can significantly enhance brain tumor classification performance from MRI images. Future research should focus on expanding dataset diversity, incorporating explainable AI frameworks, and validating the models prospectively in clinical environments to ensure safe and effective deployment in medical practice. Future studies will focus on validating the proposed model using larger and more diverse MRI datasets to ensure broader generalizability. Besides, efforts will be directed toward improving model interpretability to make the system more transparent and clinically acceptable. Optimizing the reinforcement learning component to reduce computational cost while maintaining high classification performance is another important goal. Furthermore, integrating tumor segmentation and classification into a fully automated clinical decision support system will be explored to enhance practical utility in real-world medical environments.

In future work, we plan to expand upon several directions. First, we will conduct a direct baseline comparison between the CNN-only and CNN + RL architectures trained on the same dataset to quantify the specific contribution of the reinforcement learning component. This limitation stems from the absence of a baseline CNN-only model trained on the same datasets, which would allow direct quantification of the reinforcement learning contribution. Although not feasible within the current thesis framework, this baseline comparison is prioritized for future research to further substantiate the hybrid model's advantage. Second, we intend to integrate explainable artificial intelligence (XAI) methods such as Grad-CAM or SHAP to visualize the MRI regions that influence classification outcomes, thereby improving clinical interpretability. Third, we will validate the hybrid model on larger and more diverse MRI datasets obtained from multiple

institutions to further assess its robustness and generalizability. Finally, we aim to explore the integration of our classification framework into a comprehensive clinical decision support system by combining automated tumor segmentation and classification modules. These steps will constitute the next stage of our research and are expected to strengthen both the clinical applicability and scientific impact of the proposed model.

This study was conducted as part of a Master's thesis, and due to practical constraints, several limitations remain. First, we were unable to provide a direct comparison between CNN-only and CNN+RL models trained on the same dataset. Although literature reports suggest that CNN-only models achieve high performance, our contribution lies in introducing reinforcement learning as an additional layer for adaptability and decision mapping. Second, visual interpretability methods such as Grad-CAM were not implemented. While our Q-table visualization provides insights into the agent's decision pathways, true clinical interpretability would require feature-level visualization. These aspects will be addressed in future studies by incorporating larger datasets, baseline comparisons, and advanced explainable AI methods.

In summary, while our hybrid approach shows promising results, further validation with baseline CNN-only comparisons and advanced explainability methods is necessary before clinical translation.

Authorship contribution statement

Ç. Gökçek-Saraç: Conceptualization, Supervision, Original Draft Writing, Review and Editing; **S. Arıkan:** Methodology; **Y. A. Üncü:** Conceptualization, Observation, Original Draft Writing, Review and Editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics Committee Approval and/or Informed Consent Information

As the authors of this study, we declare that we do not have any ethics committee approval and/or informed consent statement.

References

- [1] S. Rasheed, K. Rehman and M. S. H. Akash, "An insight into the risk factors of brain tumors and their therapeutic interventions", *Biomedicine & Pharmacotherapy*, 143, 112119, 2021.
- [2] Q. T. Ostrom, M. Price, C. Neff, G. Cioffi, K. A. Waite, C. Kruchko and J. S. Barnholtz-Sloan, "CBTRUS Statistical Report: Primary brain and other central nervous system tumors diagnosed in the United States in 2015-2019", *Neuro-Oncology*, 24(Suppl 5), v1-v95, 2022.
- [3] A. M. Molinaro, J. W. Taylor, J. K. Wiencke and M. R. Wrensch, "Genetic and molecular epidemiology of adult diffuse glioma", *Nature Reviews Neurology*, 15(7), 405-417, 2019.
- [4] A. Thakur, C. Faujdar, R. Sharma, S. Sharma, B. Malik, K. Nepali and J. P. Liou, "Glioblastoma: Current status, emerging targets, and recent advances", *Journal of Medicinal Chemistry*, 65(13), 8596-8685, 2022.
- [5] T. Wang, Y. Ni and L. Liu, "Innovative imaging techniques for advancing cancer diagnosis and treatment", *Cancers*, 16(14), 2607, 2024.
- [6] T. Yousaf, G. Dervenoulas and M. Politis, "Advances in MRI methodology", *International Review of Neurobiology*, 141, 31-76, 2018.
- [7] Z. Zhou and Z. R. Lu, "Gadolinium-based contrast agents for magnetic resonance cancer imaging", *Wiley Interdisciplinary Reviews: Nanomedicine and Nanobiotechnology*, 5(1), 1-18, 2013.
- [8] Y. Yu, D. H. Lee, S. L. Peng, K. Zhang, Y. Zhang, S. Jiang, et al., "Assessment of glioma response to radiotherapy using multiple MRI biomarkers with manual and semiautomated segmentation algorithms", *Journal of Neuroimaging*, 26(6), 626-634, 2016.

- [9] D. Gala, H. Behl, M. Shah and A. N. Makaryus, "The role of artificial intelligence in improving patient outcomes and future of healthcare delivery in cardiology: A narrative review of the literature", *Healthcare (Basel)*, 12(4), 481, 2024.
- [10] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed., Pearson, 2021.
- [11] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. W. M. van der Laak, B. van Ginneken and C. I. Sánchez, "A survey on deep learning in medical image analysis", *Medical Image Analysis*, 42, 60-88, 2017.
- [12] I. Hammad and K. El-Sankary, "Impact of approximate multipliers on VGG deep learning network", *IEEE Access*, 6, 60438-60444, 2018.
- [13] S. H. Wang and Y. D. Zhang, "DenseNet-201-based deep neural network with composite learning factor and precomputation for multiple sclerosis classification", *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 16(2s), 1-19, 2020.
- [14] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed., MIT Press, 2018.
- [15] L. Alzubaidi, J. Zhang, A. J. Humaidi, A. Al-Dujaili, Y. Duan, O. Al-Shamma, J. Santamaría, M. A. Fadhel, M. Al-Amidie and L. Farhan, "Review of deep learning: concepts, CNN architectures, challenges, applications, future directions", *Journal of Big Data*, 8(1), 53, 2021.
- [16] Y. N. Kuan, K. M. Goh and L. L. Lim, "Systematic review on machine learning and computer vision in precision agriculture: Applications, trends, and emerging techniques", *Engineering Applications of Artificial Intelligence*, 148, 110401, 2025.
- [17] D. Dablain, K. N. Jacobson, C. Bellinger, M. Roberts and N. M. Chawla, "Understanding CNN fragility when learning with imbalanced data", *Machine Learning*, 113, 4785-4810, 2024.
- [18] S. M. Polisetty, "AI-driven diagnosis and treatment recommendation in healthcare: A hybrid deep learning framework", *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 2025.
- [19] A. B. Sicilia, X. Zhao, A. Sosnovskikh and S. J. Hwang, "PAC Bayesian performance guarantees for deep (stochastic) networks in medical imaging", *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2021*, 12903, 560-570, 2021.
- [20] X. Jiang, Z. Hu, S. Wang and Y. Zhang, "Deep learning for medical image-based cancer diagnosis", *Cancers*, 15(14), 3608, 2023.
- [21] A. K. Mandle, S. Sahu and G. P. Gupta, "CNN-based deep learning technique for the brain tumor identification and classification in MRI images", *International Journal of Software Science and Computational Intelligence*, 14(1), 1-16, 2022.
- [22] S. S. A. Khan, A. Prova and U. Acharjee, "MRI-based brain tumor image classification using CNN", *Asian Journal of Research in Computer Science*, 15(13), 1-10, 2023.
- [23] M. A. Amou, K. Xia, S. Kamhi, and M. Mouhafid, "A Novel MRI Diagnosis Method for Brain Tumor Classification Based on CNN and Bayesian Optimization," *Healthcare*, 10 (3), 494, 2022.
- [24] J. Cheng, W. Yang, M. Huang, W. Huang, J. Jiang, Y. Zhou, R. Yang, J. Zhao, Y. Feng, Q. Feng and W. Chen, "Retrieval of brain tumors by adaptive spatial pooling and Fisher vector representation", *PloS One*, 11(6), e0157112, 2016.
- [25] M. I. Sharif, J. P. Li, M. A. Khan and M. A. Saleem, "Active deep neural network features selection for segmentation and recognition of brain tumors using MRI images", *Pattern Recognition Letters*, 129, 181-189, 2020.
- [26] R. A. Zeineldin, M. E. Karar, O. Burgert and F. Mathis-Ullrich, "Multimodal CNN networks for brain tumor segmentation in MRI: A BraTS 2022 challenge solution", *International MICCAI Brainlesion Workshop*, 13409, 127-137, 2022.
- [27] M. Attique, G. Gilanie, M. S. Mehmood, M. S. Naweed, M. Ikram, J. A. Kamran and A. Vitkin, "Colorization and automated segmentation of human T2 MR brain images for characterization of soft tissues", *PloS One*, 7(3), e33616, 2012.
- [28] A. Biswas, S. Abedin, and M. A. Kabir, "Moving Object Detection Using Ultrasonic Radar with Proper Distance, Direction, and Object Shape Analysis," *Journal of Information Systems Engineering and Business Intelligence*, 6, 2, 2020.
- [29] K. Neamah, F. Mohamed, S. R. Waheed, W. H. M. Kurdi, A. Y. Taha and K. A. Kadhim, "Utilizing deep improved ResNet50 for brain tumor classification based MRI", *IEEE Open Journal of the Computer Society*, 5, 446-456, 2024.
- [30] H. A. Munira and M. S. Islam, "Hybrid deep learning models for multi-classification of tumour from brain MRI", *Journal of Information Systems Engineering and Business Intelligence*, 8(2), 162-174, 2022.
- [31] P. Pilaoon, A. Narkthewan, N. Wadlom, R. Varakulsiripunth, K. Hamamoto and N. Maneerat, "Glioma brain tumor classification using transfer learning", *Proceedings of the 2024 10th International Conference on Engineering, Applied Sciences, and Technology (ICEAST)*, 17-22, 2024.

- [32] P. Pilagoon, A. Narkthewan, N. Wadlom, R. Varakulsiripunth, K. Hamamoto and N. Maneerat, "Glioma brain tumor classification using transfer learning", Proceedings of the 10th International Conference on Engineering, Applied Sciences and Technology (ICEAST), 17–22, 2024.
- [33] M. A. Khan, M. Z. Hussain, S. Mehmood, M. F. Khan, M. Ahmad, T. Mazhar, T. Shahzad, and M. M. Saeed, "Transfer learning for accurate brain tumor classification in MRI: A step forward in medical diagnostics," *Discover Oncology*, vol. 16, no. 1040, 2025.
- [34] A. Sevinc, M. Uçan and B. Kaya, "A distillation approach to transformer-based medical image classification with limited data", *Diagnostics*, 15(7), 929, 2025.