Minimax Theory and its Applications

Volume 11 (2026), No. 3s, 84 to 107

Quality-Conditioned Region-Specific Learning for Multimodal 3D Brain Tumor Segmentation

Felicia Rajathi P[1], Dr. D.Gladis[2]

[1]Department of Computer Science, Madras Christian College (Autonomous), Chennai, Tamil Nadu, India
Email: felicia@mcc.edu.in
[2]Bharathi Women's College (Autonomous), Chennai – 600108, Tamil Nadu, India,
Email: gladischristopher@gmail.com

Abstract
Accurate segmentation of brain tumors from multimodal Magnetic Resonance Imaging (MRI) is essential for quantitative tumor assessment and subsequent clinical analysis. However, segmentation of the enhancing tumor (ET) region remains particularly challenging because of its relatively small and heterogeneous appearance and its sensitivity to variations in MRI image quality. Although existing deep learning approaches have achieved substantial progress in multimodal brain tumor segmentation, most methods primarily focus on improving network architectures or loss functions, while the influence of modality-specific MRI quality on region-specific segmentation learning remains insufficiently addressed. In particular, the integration of quantitative MRI quality information with explicit learning of the distinct whole tumor (WT), tumor core (TC), and ET regions has received limited investigation. To address this gap, this study proposes a Quality-Conditioned Region-Specific Learning (QCR) framework for multimodal 3D brain tumor segmentation. The framework incorporates quantitative MRI quality characteristics into the segmentation network and explicitly conditions the learning process on WT, TC, and ET regions. Four MRI modalities, T1, T1ce, T2, and FLAIR, from the BraTS 2020 dataset were used. A 16-dimensional quality representation was constructed using coefficient of variation (CVP), signal-to-noise ratio measures (SNR3 and SNR4), and coefficient of joint variation (CJV) across the four modalities. The experiments included 368 valid subjects, divided into 257 training, 55 validation, and 56 independent testing subjects. The proposed QCR framework was implemented using a controlled 3D U-Net with Dice and weighted cross-entropy loss. On the independent test set, QCR achieved mean Dice scores of 0.8888 ± 0.1067 for WT, 0.7973 ± 0.2500 for TC, and 0.7331 ± 0.2360 for ET. Notably, the ET Dice improved from 0.7075 in the controlled baseline to 0.7331, representing an absolute improvement of 0.0256. The corresponding ET IoU, precision, and sensitivity were 0.6222 ± 0.2427, 0.7608 ± 0.2259, and 0.8084 ± 0.2216, respectively. An ablation study further showed that QCR produced comparable region-specific Dice performance to fixed-region conditioning, while providing a quality-conditioned formulation that incorporates MRI quality characteristics into the learning process. These findings demonstrate the potential of combining quantitative MRI quality information with region-specific learning, particularly for improving segmentation of the challenging ET region.

Keywords: Brain tumor segmentation; Enhancing tumor; Multimodal MRI; Quality-conditioned learning; Region-specific learning; 3D U-Net; BraTS 2020; MRI quality assessment.

1. Introduction

Brain tumors are among the challenging abnormalities encountered in medical image analysis because of their heterogeneous appearance, irregular boundaries, and variation in tissue characteristics. Magnetic resonance imaging (MRI) is widely used for brain tumor assessment because different MRI sequences provide complementary information about tumor tissue and surrounding structures. In particular, multimodal MRI sequences such as T1, contrast-enhanced T1 (T1ce), T2, and FLAIR provide complementary anatomical and pathological information that can support automated tumor delineation. Recent reviews have highlighted the increasing use of deep learning for multimodal brain tumor segmentation and the continued development of convolutional neural networks, U-Net variants, Transformer-based models, and hybrid architectures [1,2].

Automated brain tumor segmentation is important for obtaining quantitative measurements that can support diagnosis, treatment planning, radiotherapy targeting, and longitudinal assessment [2,3]. The BraTS benchmark

has become an important reference dataset for evaluating automated glioma segmentation methods and commonly evaluates three clinically relevant tumor regions: whole tumor (WT), tumor core (TC), and enhancing tumor (ET) [4]. Recent reviews of the BraTS datasets have also emphasized their continuing role in developing and evaluating automated MRI segmentation methods [4].

Despite considerable progress, accurate segmentation of individual tumor subregions remains challenging. The enhancing tumor (ET) region is of particular interest because it can exhibit heterogeneous intensity patterns and relatively complex boundaries. Recent literature continues to report the difficulty of achieving consistent performance across tumor subregions, particularly for ET and TC [2]. Consequently, recent methods have investigated multi-scale feature extraction, attention mechanisms, feature fusion, and improved encoder-decoder architectures to improve the representation of heterogeneous tumor structures [5,6]. For example, recent multimodal segmentation work has incorporated multi-scale contextual aggregation and gated attention fusion to improve feature representation and suppress irrelevant information in BraTS 2020 MRI data [5].

The continued development of architectural modules demonstrates the importance of learning discriminative spatial and contextual features. However, the reliability of the input MRI information itself is another factor that can influence segmentation performance. Multimodal MRI images can exhibit differences in intensity characteristics, noise levels, contrast, and quality across modalities and subjects. Recent research has increasingly recognized robustness to degraded, missing, or variable-quality MRI inputs as an important issue for reliable multimodal brain tumor segmentation [2,7]. Nevertheless, most segmentation frameworks primarily address these challenges through architectural modifications, multimodal fusion strategies, attention mechanisms, or robust-learning techniques rather than explicitly incorporating a quantitative representation of MRI quality into region-specific segmentation learning.

This leads to an important research gap. Existing multimodal segmentation approaches generally exploit the complementary information contained in multiple MRI modalities, but the quantitative quality characteristics of these modalities are not commonly used as an explicit conditioning signal for learning different tumor regions. Moreover, WT, TC, and ET have different spatial and morphological characteristics, meaning that a common segmentation strategy may not adequately account for the specific difficulties associated with each region. Recent research has identified persistent subregion-specific failures, particularly in ET and TC, and has emphasized the need for more robust approaches that consider input reliability and region-specific behavior [2].

To address this gap, this study proposes Quality-Conditioned Region-Specific Learning (QCR) for multimodal 3D brain tumor segmentation. The proposed framework incorporates quantitative MRI quality characteristics into the segmentation learning process while explicitly considering the distinct WT, TC, and ET regions. A 16-dimensional MRI quality representation is constructed using coefficient of variation (CVP), signal-to-noise ratio measures (SNR3 and SNR4), and coefficient of joint variation (CJV) across T1, T1ce, T2, and FLAIR modalities. The framework is evaluated using a controlled 3D U-Net architecture with a hybrid Dice and weighted cross-entropy loss on the BraTS 2020 dataset.

The principal contributions of this study are as follows:

1. **Quality-conditioned learning:** A quantitative MRI quality representation is incorporated into the multimodal 3D segmentation framework rather than relying solely on image intensity and learned feature representations.
2. **Region-specific learning:** The framework explicitly considers the WT, TC, and ET regions, with particular attention to the challenging ET segmentation problem.
3. **Quantitative MRI quality representation:** A 16-dimensional quality vector derived from CVP, SNR3, SNR4, and CJV is used to represent quality characteristics across the four MRI modalities.
4. **Controlled evaluation and ablation analysis:** Controlled experiments and a fixed-region ablation are conducted to examine the contribution of the proposed quality-conditioned and region-specific learning strategy.
5. **ET-focused evaluation:** The proposed framework is evaluated separately for WT, TC, and ET. On the independent test set, QCR achieved Dice scores of 0.8888 ± 0.1067, 0.7973 ± 0.2500, and 0.7331 ± 0.2360 for WT, TC, and ET, respectively. In particular, the ET Dice increased from 0.7075 for the controlled baseline to 0.7331 with QCR, providing an absolute improvement of 0.0256.

Therefore, this study investigates a different direction from approaches that rely primarily on increasingly complex segmentation architectures: whether quantitative MRI quality information can be explicitly incorporated into region-specific learning to improve multimodal 3D brain tumor segmentation, particularly for the challenging ET region.

2. Related Work

2.1 U-Net-Based Brain Tumor Segmentation

U-Net and its variants remain widely used for automated brain tumor segmentation because the encoder-decoder architecture can combine high-level contextual information with spatial information through skip connections. A recent review of 35 U-Net-based brain tumor segmentation studies published between 2019 and 2025 reported continued development from conventional 2D and 3D U-Net architectures toward models incorporating residual blocks, attention mechanisms, dense connectivity, and Transformer components [8]. The review also identified BraTS as an important benchmark for evaluating segmentation of tumor subregions and noted continuing challenges associated with small and irregular tumor regions and model generalization.

Recent work has further enhanced U-Net through stronger encoders and multi-scale feature extraction. Preetha *et al.* proposed a multi-scale attention U-Net with an EfficientNetB4 encoder for brain tumor segmentation, using multi-resolution representations and attention mechanisms to improve tumor feature extraction [9]. These approaches demonstrate the continued effectiveness of U-Net-based architectures, but their primary emphasis is on architectural feature extraction, rather than explicitly incorporating quantitative MRI-quality characteristics into segmentation learning.

2.2 Attention and Multi-Scale Feature Learning

Attention mechanisms have become an important research direction for brain tumor segmentation because they can emphasize informative tumor features while reducing irrelevant background information. Aslam *et al.* proposed a multimodal multi-scale contextual aggregation and attention-fusion framework using T1, T2, FLAIR, and T1-CE MRI from BraTS 2020 [10]. Their approach combines multi-scale contextual aggregation with gated attention fusion to improve representation of heterogeneous tumor structures and suppress irrelevant information.

Similarly, Preetha *et al.* incorporated multi-scale attention into U-Net together with an EfficientNetB4 encoder to improve feature representation at different resolutions [9]. Yang *et al.* proposed HLNet, a U-Net++-based architecture incorporating a high-level attention mechanism for MRI brain tumor segmentation [11].

These studies demonstrate that attention and multi-scale feature learning can improve the extraction of tumor-related information. However, these approaches primarily determine feature importance from the image-derived representations learned by the network. The QCR framework follows a different direction by investigating whether quantitative MRI quality information can explicitly condition the segmentation learning process.

2.3 Multimodal MRI and Feature Fusion

Multimodal MRI segmentation benefits from the complementary information provided by different sequences, including T1, T1ce, T2, and FLAIR. Consequently, multimodal feature fusion has become a major component of modern brain tumor segmentation frameworks. A 2025 review of multimodal brain tumor segmentation analyzed U-Net, Transformer, hybrid, and emerging Mamba-based approaches and emphasized multimodal fusion, missing-modality handling, and dataset heterogeneity as important research areas [12].

The review also noted that existing fusion strategies may not dynamically adapt the contribution of different modalities to individual cases [12]. This observation is relevant to the present study because QCR does not treat all available modality information as equally reliable. Instead, the proposed framework introduces a quantitative representation of MRI quality to provide additional information about the characteristics of the input modalities.

2.4 Transformer and Emerging Architectures

Recent brain tumor segmentation research has expanded beyond conventional CNNs and U-Net variants toward Transformer, hybrid CNN-Transformer, Mamba, diffusion, and other emerging architectures. Verma and Yadav reviewed CNN-based, U-Net-based, transfer-learning-based, Transformer-based, and hybrid Transformer approaches and identified robustness, multimodal MRI, domain adaptation, class imbalance, and generalization as continuing challenges [13].

The 2025 multimodal segmentation review similarly reported increasing use of Transformer and Mamba-based U-Net architectures while emphasizing limitations related to multimodal fusion, missing modalities, heterogeneous datasets, and generalization [12].

More recently, Zhou *et al.* provided a 2026 deployment-oriented review of multimodal brain tumor segmentation and emphasized that high benchmark accuracy does not necessarily guarantee clinical reliability. Their analysis identifies input reliability, degraded or missing modalities, cross-center shifts, and subregion-specific failures, particularly in ET and TC, as persistent challenges [14].

These findings suggest that continued increases in architectural complexity alone may not fully address the reliability problem. The present study therefore investigates a complementary strategy based on quantitative input-quality conditioning.

2.5 Enhancing Tumor and Region-Specific Segmentation

Brain tumor segmentation commonly evaluates the whole tumor (WT), tumor core (TC), and enhancing tumor (ET). Although WT generally contains a larger and more spatially extensive region, ET can be considerably more difficult to delineate because of its smaller, heterogeneous, and irregular characteristics. The recent U-Net review specifically identifies small and irregularly bounded subregions such as ET as an ongoing research gap [8].

Recent methods have attempted to improve difficult subregion segmentation through multi-scale contextual learning and attention mechanisms. Aslam *et al.* reported the use of multi-scale contextual aggregation and gated attention fusion to improve representation of heterogeneous tumor regions [10]. Similarly, multi-scale attention U-Net approaches have been developed to improve feature representation across different spatial resolutions [9].

However, these approaches mainly improve the representation of tumor features, whereas the present study additionally considers the quality characteristics of the MRI inputs. This distinction is particularly relevant to ET, where small segmentation errors can have a substantial effect on overlap-based evaluation.

2.6 MRI Quality, Robustness, and Input Reliability

Recent research has increasingly recognized that MRI variability and input reliability can affect segmentation performance. Variations in acquisition protocols, scanners, image quality, missing modalities, and cross-center distributions can create conditions in which a model that performs well on a standardized benchmark may not maintain the same reliability under different inputs [12–14].

The 2026 survey by Zhou *et al.* particularly emphasizes a reliability chain connecting input quality, multimodal fusion, architecture, and subregion-specific failure. It identifies persistent ET/TC fragility under missing, degraded, and shifted MRI conditions [14].

However, input reliability and quantitative MRI-quality conditioning are not identical concepts. Robust learning and missing-modality methods generally attempt to maintain segmentation performance when information is unavailable, degraded, or shifted. In contrast, the QCR framework considers the situation in which the MRI modalities are available but exhibit measurable differences in their quality characteristics. The proposed approach therefore represents MRI quality explicitly using CVP, SNR3, CJV, and SNR4 and investigates its use as a conditioning signal for segmentation.

2.7 Research Gap

The recent literature demonstrates substantial progress in U-Net-based architectures, attention and multi-scale feature learning, multimodal fusion, Transformer/Mamba architectures, and robust learning under missing or

degraded modalities [8–14]. These developments have substantially improved automated brain tumor segmentation; nevertheless, several issues remain relevant to reliable WT, TC, and particularly ET segmentation.

The following observations motivate the present study:

1. **ET remains a persistent difficult subregion.**

   Despite improvements in overall segmentation performance, accurate delineation of ET remains challenging because of its relatively small, heterogeneous, and irregular appearance. Recent reviews continue to identify small and irregular subregions, particularly ET, as an unresolved problem [8], while the latest multimodal segmentation survey reports persistent ET/TC fragility under degraded, missing, and shifted MRI conditions [14].

2. **Most architectural improvements focus on feature extraction, attention, fusion, or backbone design.**
   Recent studies have introduced multi-scale attention [9], contextual aggregation and attention fusion [10], high-level attention mechanisms [11], Transformer-based architectures, and emerging Mamba-based models [12,13]. Although these approaches improve feature representation, they generally focus on **how the network extracts and combines information from the images**, rather than explicitly conditioning segmentation on quantitative MRI-quality characteristics.

3. **Input reliability and region-specific segmentation are generally investigated as related but separate issues.**

   Recent literature increasingly recognizes missing modalities, degraded MRI, cross-center variation, and input reliability as important challenges [12,14]. However, the availability of an MRI modality does not necessarily imply that its information has equivalent quality or reliability. This creates an opportunity to investigate whether **quantitative MRI-quality information can be incorporated explicitly into segmentation learning together with region-specific learning for WT, TC, and ET**.

4. **The relationship between quantitative MRI quality and region-specific segmentation remains insufficiently explored.**

   Existing approaches generally learn segmentation features directly from MRI data and optimize segmentation losses. In contrast, the present study introduces a **16-dimensional MRI quality representation** derived from **CVP, SNR3, CJV, and SNR4** across T1, T1ce, T2, and FLAIR. This allows measurable MRI-quality characteristics to be explicitly represented rather than leaving input-quality variation implicit.

5. **ET should be explicitly considered when evaluating a quality-conditioned segmentation framework.**
   In the present experiments, the controlled baseline obtained an ET Dice of **0.7075**, whereas the QCR framework achieved **0.7331 ± 0.2360** on the independent test set, corresponding to an absolute improvement of **0.0256**. Therefore, ET performance is considered explicitly in the evaluation rather than relying only on overall or WT performance.

Accordingly, the present study investigates Quality-Conditioned Region-Specific Learning (QCR), in which quantitative MRI-quality characteristics are explicitly represented and incorporated into a 3D segmentation framework together with region-specific learning. The objective is not simply to introduce another attention, multi-scale, frequency, or Transformer module, since these directions are already well represented in recent literature. Instead, this study investigates a quality-aware learning perspective, in which characteristics of the input MRI data are considered alongside the target tumor regions.

The central research question addressed by this study is:

Can quantitative MRI-quality information, when incorporated into region-specific learning, improve multimodal 3D brain tumor segmentation, particularly for the challenging enhancing tumor (ET) region?

This research question forms the basis for the proposed QCR framework and the subsequent controlled and ablation experiments.

3. Proposed Quality-Conditioned Region-Specific Learning Framework

3.1 Overview

The proposed Quality-Conditioned Region-Specific Learning (QCR) framework is designed to investigate whether quantitative MRI quality information can be incorporated into multimodal 3D brain tumor segmentation to support region-specific learning. The framework retains the four MRI modalities—T1, T1ce, T2, and FLAIR—as the primary image inputs. The MRI quality information does not replace the original MRI images; instead, it provides an additional conditioning signal to the segmentation model.

The overall QCR framework consists of two complementary information pathways. The first pathway processes the four MRI modalities through the 3D U-Net segmentation network to learn volumetric tumor representations. The second pathway processes a 16-dimensional MRI quality vector obtained from quantitative image-quality measurements. The quality vector is passed through a quality-conditioning multilayer perceptron (MLP), producing three region-specific outputs corresponding to WT, TC, and ET. These outputs are incorporated into the region-specific learning objective.

3.2 Multimodal MRI Input

Four multimodal MRI sequences from the BraTS 2020 dataset are used as the input:

- T1
- T1ce
- T2
- FLAIR

The four modalities are provided to the 3D U-Net as a four-channel volumetric input. The segmentation target contains the tumor regions required for evaluation, namely whole tumor (WT), tumor core (TC), and enhancing tumor (ET).

The QCR framework therefore preserves the information available from the original multimodal MRI data while providing an additional quality-related representation.

3.3 Statistical Selection of MRI Quality Features

To construct a meaningful quality representation for the proposed Quality-Conditioned Region-Specific (QCR) framework, a statistical screening procedure was performed before defining the final quality vector. The objective was to identify complementary MRI-quality measures while reducing redundancy among highly related metrics. The analysis considered quantitative image-quality measures obtained from the four MRI modalities, namely T1, T1ce, T2, and FLAIR.

The initial set of MRI quality measures comprised CV, PSNR, SNR1, SNR2, SNR3, SNR4, CNR, CVP, CJV, and EFC. Measures such as CPP and FBER, together with metadata and embedding-coordinate variables, were excluded from the quality representation because they were not retained as suitable complementary quality descriptors in the screening process.

The measures were examined according to their variability, distribution, redundancy, and representation of different quality characteristics across the four MRI modalities. Spearman rank correlation analysis was used to investigate inter-metric relationships and identify highly redundant measures. This step was important because including several strongly correlated quality measures could increase the dimensionality of the quality representation without providing substantially different information.

The measures were also considered according to their principal quality characteristics. CV and CJV represented intensity variation and inhomogeneity, while SNR1–SNR4 represented signal-related characteristics. CNR represented contrast-related information, CVP represented spatial/intensity variation, and EFC represented an

information or entropy-related characteristic. This analysis was used to retain quality measures representing complementary rather than substantially overlapping information.

In addition, a patient-level validation analysis was performed using the 55 validation subjects to examine the relationship between MRI-quality measures and segmentation performance. Spearman associations were calculated between the quality measures and the Dice scores of the WT, TC, and ET regions. The analysis provided additional evidence regarding the relevance of the quality measures to region-specific segmentation behaviour.

Based on the combined screening, redundancy analysis, quality-domain coverage, and validation-level relevance analysis, four complementary quality measures were selected for the QCR representation: CVP, SNR3, CJV, and SNR4. Each selected measure was calculated separately for T1, T1ce, T2, and FLAIR. Consequently, the final quality representation contained 16 quality features.

3.4 Final 16-Dimensional MRI Quality Representation

A central component of QCR is the representation of MRI quality using quantitative image-quality measurements. The quality measurements were obtained using **MRQy**, which provides quantitative image-quality metrics for the MRI modalities.

Following the metric-selection procedure described in Section 3.3, the final QCR representation contains four quality-measurement groups:

- **CVP** for T1, T1ce, T2, and FLAIR
- **SNR3** for T1, T1ce, T2, and FLAIR
- **CJV** for T1, T1ce, T2, and FLAIR
- **SNR4** for T1, T1ce, T2, and FLAIR

Thus, four metrics × four MRI modalities produce a 16-dimensional quality vector:

q = [CVP_T1, CVP_T1ce, CVP_T2, CVP_FLAIR, SNR3_T1, SNR3_T1ce, SNR3_T2, SNR3_FLAIR, CJV_T1, CJV_T1ce, CJV_T2, CJV_FLAIR, SNR4_T1, SNR4_T1ce, SNR4_T2, SNR4_FLAIR] (1)

The quality representation is used as an additional conditioning signal rather than as a replacement for the MRI images.

To prevent information leakage, missing quality values were imputed using training-set medians, and z-score normalization parameters were calculated using the training data only. The same parameters were subsequently applied unchanged to the validation and test sets.

The resulting 16-dimensional quality representation is provided as input to the quality-conditioning network described in Section 3.5.

3.5 Quality-Conditioning Network

The selected 16-dimensional MRI quality vector is provided to a lightweight quality-conditioning network to generate region-specific quality weights for the three tumor regions: Whole Tumor (WT), Tumor Core (TC), and Enhancing Tumor (ET).

The quality-conditioning network is implemented as a multilayer perceptron (MLP) with the following architecture:16 → 32 → 16 → 3

The first fully connected layer maps the 16-dimensional quality representation to 32 hidden features. A ReLU activation is then applied to introduce non-linearity. A dropout rate of 0.1 is used for regularization. The second fully connected layer reduces the representation from 32 to 16 features, followed by another ReLU activation. Finally, a fully connected layer maps the 16-dimensional representation to three outputs corresponding to the WT, TC, and ET regions.

The three outputs are normalized using the Softmax function to obtain region-specific quality weights:

w = [w_WT, w_TC, w_ET](2)

where w_WT, w_TC, and w_ET represent the quality-conditioned weights associated with Whole Tumor, Tumor Core, and Enhancing Tumor, respectively.

The Softmax operation is expressed as:

w_r = exp(z_r) / Σ exp(z_k)(3)

where z_r denotes the output of the final fully connected layer for region r, and the summation is performed over the three tumor regions.

These quality-derived weights are subsequently incorporated into the region-specific learning component of the proposed QCR framework. Thus, the network can assign different levels of emphasis to WT, TC, and ET according to the quantitative MRI-quality characteristics of each subject.

The quality-conditioning network therefore provides a mechanism for transforming subject-specific MRI-quality information into region-specific learning weights, while the original multimodal MRI data remain the primary input to the 3D segmentation network.

3.6 Region-Specific Learning

The proposed QCR framework incorporates the quality information obtained from the MRI-quality representation into region-specific learning for the three clinically relevant tumor regions: Whole Tumor (WT), Tumor Core (TC), and Enhancing Tumor (ET).

Let the quality-conditioning network produce the three normalized weights:

w = [w_WT, w_TC, w_ET](4)

These weights provide a subject-specific representation of the relative quality conditions associated with the three tumor regions. Instead of applying the same learning emphasis to all tumor regions, the QCR framework uses these quality-derived weights to modulate the regional segmentation objective.

For a given subject, the region-specific loss can be represented as:

L_Region = w_WT L_WT + w_TC L_TC + w_ET L_ET(5)

where L_WT, L_TC, and L_ET denote the segmentation losses associated with the Whole Tumor, Tumor Core, and Enhancing Tumor regions, respectively.

This formulation enables the segmentation model to incorporate subject-specific MRI-quality information into regional learning. In particular, the approach provides additional emphasis to regions according to the quality characteristics represented by the 16-dimensional MRI-quality vector.

The region-specific component is combined with the conventional hybrid segmentation objective. The final QCR training objective is therefore defined as:

L_QCR = L_Hybrid + 0.2 L_Region(6)

where 0.2 is the weighting coefficient assigned to the region-specific learning component.

The proposed formulation maintains the multimodal MRI information as the primary segmentation input while introducing quantitative image-quality information as an additional conditioning signal. This allows the model to learn quality-conditioned regional segmentation without replacing the underlying 3D U-Net segmentation architecture.

3.7 Hybrid Segmentation Loss

The proposed QCR framework uses a hybrid segmentation loss combining Dice Loss and Weighted Cross-Entropy (CE) Loss. The combination is intended to optimize both region overlap and voxel-level classification during multimodal 3D brain tumor segmentation.

The hybrid segmentation loss is defined as:

L_Hybrid = L_Dice + L_WeightedCE(7)

The Dice Loss measures the overlap between the predicted segmentation and the corresponding ground-truth region, while the weighted cross-entropy component provides voxel-level classification supervision. Class weighting is used to account for the different contributions of the background and tumor regions during training.

The class weights used in the experiments are:[0.1, 0.3, 0.3, 0.3]

where the four values correspond to the background, WT, TC, and ET classes, respectively.

The QCR framework further incorporates the region-specific learning component described in Section 3.6. Therefore, the complete training objective is formulated as:

where L_Region represents the quality-conditioned region-specific loss and 0.2 controls its contribution to the overall optimization objective.

Thus, the complete QCR objective combines conventional segmentation supervision with quality-conditioned regional learning. Binary Cross-Entropy (BCE) loss is not used in the proposed formulation.

This loss formulation was kept consistent with the controlled baseline and fixed-region ablation experiments, allowing the contribution of quality-conditioned regional learning to be examined under a controlled experimental setting.

3.8 Overall QCR Architecture

The overall architecture of the proposed Quality-Conditioned Region-Specific (QCR) framework integrates multimodal MRI segmentation with quantitative MRI-quality information. The framework uses the four MRI modalities—T1, T1ce, T2, and FLAIR—as the primary segmentation input, while the selected MRI-quality measures provide an additional conditioning signal.

The overall architecture and learning process of the proposed QCR framework are illustrated in Figure1.

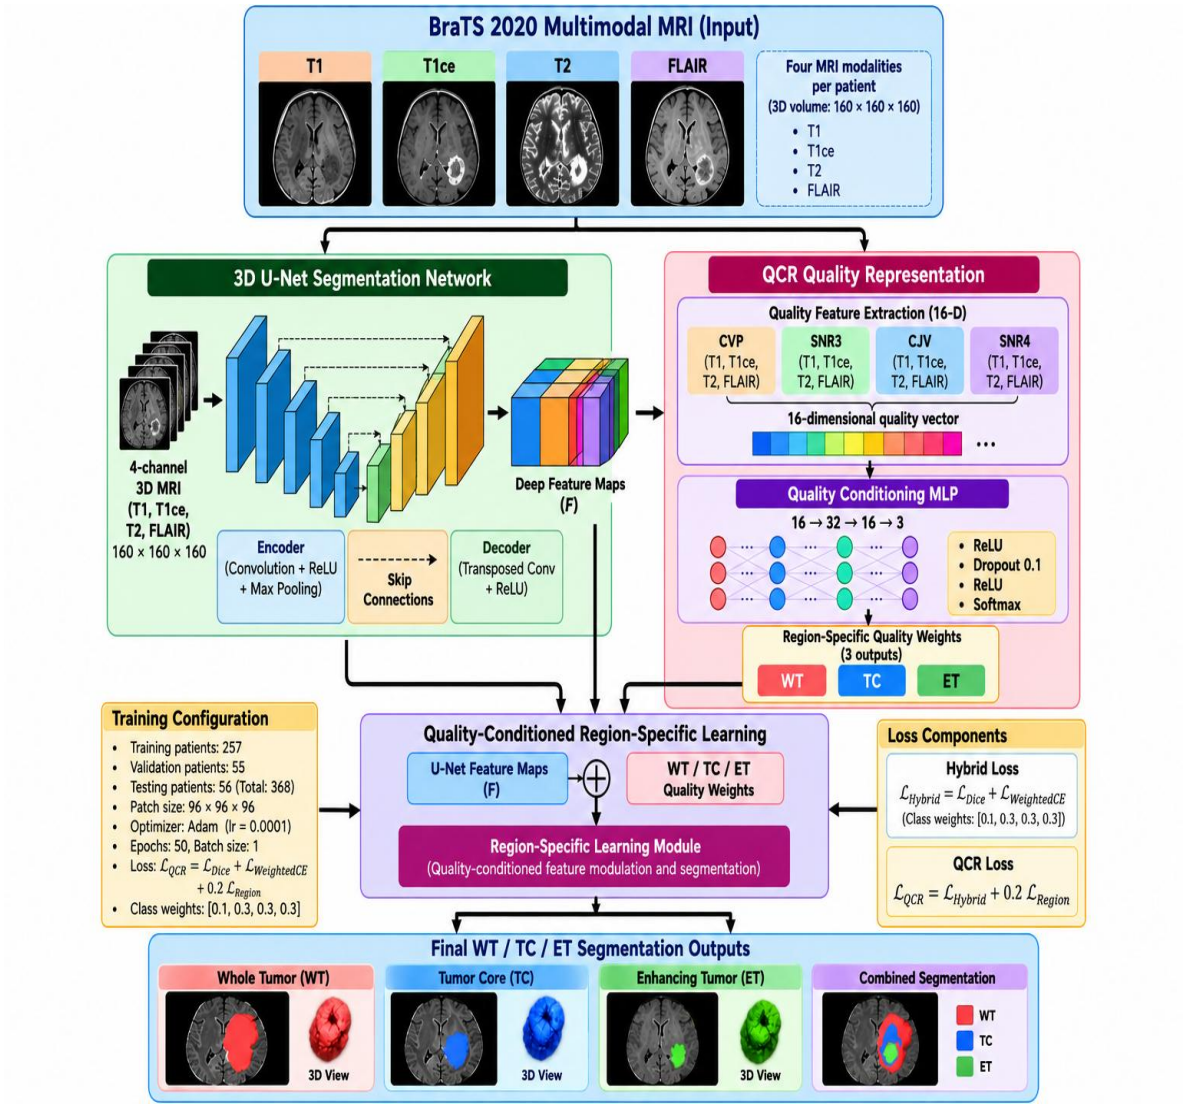


Fig. 1. Overall architecture of the proposed Quality-Conditioned Region-Specific Learning (QCR) framework for multimodal 3D brain tumor segmentation.

The quality-derived regional weights are then incorporated into the region-specific learning component of the segmentation framework. The resulting quality-conditioned regional objective is combined with the hybrid Dice and weighted cross-entropy segmentation loss.

The framework produces the final segmentation of the three tumor regions: Whole Tumor (WT), Tumor Core (TC), and Enhancing Tumor (ET).

Thus, the proposed architecture does not replace the multimodal MRI input with quality information. Instead, quantitative MRI-quality information is introduced as an additional conditioning signal that guides region-specific learning within the 3D segmentation framework. This design enables the model to consider both the anatomical information contained in multimodal MRI and subject-specific image-quality characteristics during segmentation.

3.9 Training Configuration

The QCR model was trained using the same controlled experimental protocol as the baseline to enable a fair comparison. The training and evaluation configuration used for the proposed QCR model is summarized in Table 1.

Table 1. Training Configuration of the Proposed QCR Model

Parameter	**Configuration**
Dataset	BraTS 2020
Valid subjects	368
Training	257
Validation	55
Test	56
MRI modalities	T1, T1ce, T2, FLAIR
Training patch	96 × 96 × 96
Test inference	Full-volume sliding-window
Optimizer	Adam
Learning rate	0.0001
Batch size	1
Epochs	50
Base loss	Dice + Weighted CE
Class weights	[0.1, 0.3, 0.3, 0.3]
QCR λ	0.2
Best epoch	50
Best recorded training loss	0.38597

The controlled baseline and QCR models used the same dataset partition and principal training configuration, allowing differences in segmentation performance to be examined under a controlled experimental setting. The QCR-specific region-conditioning component was incorporated through the additional region-specific loss term weighted by $\boldsymbol{\lambda = 0.2}$.

4. Experimental Setup and Methodology

4.1 Dataset

The experiments were conducted using the BraTS 2020 multimodal brain tumor MRI dataset. Each subject contains four MRI modalities: T1-weighted (T1), contrast-enhanced T1-weighted (T1ce), T2-weighted (T2), and Fluid-Attenuated Inversion Recovery (FLAIR) images. These complementary modalities provide different representations of tumor anatomy and pathology.

After excluding the invalid subject BraTS20_Training_355, a total of 368 valid subjects were included in the experiments. The dataset was divided into training, validation, and testing subsets containing 257, 55, and 56 subjects, respectively. A fixed split with random seed 42 was used for the controlled experiments. The distribution of the valid subjects across the training, validation, and testing sets is summarized in Table 2.

Table 2. Dataset Distribution for the Experimental Study

Dataset component	**Number of subjects**
Total valid subjects	368
Training	257
Validation	55
Testing	56

The test set was kept completely separate from model training and validation and was used only for the final patient-level evaluation.

4.2 Data Preparation and Input Representation

The four MRI modalities were processed as a multimodal volumetric input, with T1, T1ce, T2, and FLAIR retained as separate input channels. For computationally controlled training, the segmentation network operated on 96 × 96 × 96 voxel patches.

The same data partitioning and evaluation protocol were maintained across the controlled baseline, fixed-region ablation, and QCR experiments to provide a consistent basis for comparison.

For final testing, predictions were generated using full-volume sliding-window inference, and patient-level segmentation metrics were subsequently calculated.

4.3 Tumor Regions

The segmentation task evaluated three tumor regions:

- **Whole Tumor (WT):** the complete tumor region.
- **Tumor Core (TC):** the tumor core excluding surrounding edema.
- **Enhancing Tumor (ET):** the enhancing component of the tumor.

Particular attention was given to the ET region because it represents a challenging segmentation target and was also the lowest-performing region in the controlled baseline.

4.4 Baseline Segmentation Network

A controlled **3D U-Net** was used as the primary segmentation architecture. The same underlying segmentation architecture was retained for the controlled baseline, fixed-region ablation, and QCR experiments so that the effect of region-specific and quality-conditioned learning could be examined without introducing a different backbone architecture.

The network receives the four-channel multimodal MRI volume and generates voxel-level tumor segmentation predictions. Three-dimensional encoder–decoder processing is used to learn volumetric representations, with skip connections transferring spatial information from the encoder to the corresponding decoder stages.

4.5 Training Configuration

The networks were trained using the **Adam optimizer** with a learning rate of $\mathbf{1 \times 10^{-4}}$ and a batch size of **1** for **50 epochs**. The base segmentation objective combined Dice loss with weighted cross-entropy loss using class weights of **[0.1, 0.3, 0.3, 0.3]**.

For the QCR model, an additional region-specific loss term was incorporated with a weighting coefficient of **0.2**.The training configuration used for the controlled 3D U-Net, fixed-region ablation, and QCR model is summarized in Table 3.

Table 3. Training Configuration for the Controlled and QCR Models

Parameter	**Configuration**
Architecture	3D U-Net
MRI modalities	T1, T1ce, T2, FLAIR
Training patch	96 × 96 × 96
Optimizer	Adam
Learning rate	1×10^{-4}
Batch size	1
Training epochs	50
Base loss	Dice + Weighted Cross-Entropy
Class weights	[0.1, 0.3, 0.3, 0.3]

Parameter	Configuration
QCR region-loss weight	0.2
Random seed	42

For the QCR experiment, the best recorded training epoch was **epoch 50**, with a training loss of **0.38597**. The corresponding fixed-region ablation reached its best recorded training loss of **0.40308** at epoch 50.

4.6 Evaluation Protocol

The final models were evaluated on the 56-patient held-out test set. Full-volume predictions were obtained using sliding-window inference, and patient-level metrics were calculated independently for WT, TC, and ET.

The following metrics were used:

Dice coefficient

$$Dice = \frac{2TP}{2TP+FP+FN} \quad (8)$$

Intersection over Union

$$IoU = \frac{TP}{TP+FP+FN} \quad (9)$$

Precision

$$Precision = \frac{TP}{TP+FP} \quad (10)$$

Sensitivity

$$Sensitivity = \frac{TP}{TP+FN} \quad (11)$$

Specificity

$$Specificity = \frac{TN}{TN+FP} \quad (12)$$

The results are reported as mean ± standard deviation across the 56 test subjects.

4.7 Controlled Experimental Design

Three experimental configurations were evaluated:

1. **Controlled 3D U-Net baseline:** Standard multimodal segmentation without region-specific QCR conditioning.
2. **Fixed-region 3D U-Net:** Region-specific learning using fixed regional weighting.
3. **QCR-3D U-Net:** Quality-conditioned region-specific learning using patient-specific MRI-quality information.

The controlled baseline achieved Dice scores of 0.8355, 0.7453, and 0.7075 for WT, TC, and ET, respectively. The QCR model achieved 0.8888, 0.7973, and 0.7331, corresponding to absolute improvements of 0.0533, 0.0520, and 0.0256, respectively.

The fixed-region ablation achieved Dice scores of 0.8933, 0.8002, and 0.7331 for WT, TC, and ET, respectively. Paired Wilcoxon testing between the fixed-region and QCR models showed no statistically significant difference in Dice for WT ($p = 0.1511$), TC ($p = 0.7941$), or ET ($p = 0.7208$).

Therefore, the experimental design allows the effect of fixed region-specific learning to be compared with the quality-conditioned formulation, while the results indicate that region-specific learning accounts for the principal improvement over the controlled baseline. The results indicate that region-specific learning accounts for the principal improvement over the controlled baseline, while QCR provides a patient-specific quality-conditioned extension with comparable Dice performance to fixed regional weighting. QCR also exhibited a consistent shift toward higher sensitivity and lower precision relative to the fixed-region model.

5. Experimental Results and Analysis

5.1 Overall Segmentation Performance

The segmentation performance of the controlled 3D U-Net, fixed-region ablation, and proposed QCR model was evaluated on the same 56-patient held-out test set. The comparison focuses on Dice coefficient for the three tumor regions: Whole Tumor (WT), Tumor Core (TC), and Enhancing Tumor (ET). The Dice performance of the three experimental configurations is summarized in Table 4.

Table 4. Comparison of Dice performance across the three experimental configurations

Model	WT Dice	TC Dice	ET Dice
Controlled 3D U-Net	0.8355	0.7453	0.7075
Fixed-Region 3D U-Net	0.8933	0.8002	0.7331
QCR-3D U-Net	0.8888	0.7973	0.7331

The controlled 3D U-Net achieved Dice scores of 0.8355 for WT, 0.7453 for TC, and 0.7075 for ET. Incorporating fixed region-specific learning increased the Dice scores to 0.8933, 0.8002, and 0.7331, respectively.

The proposed QCR model achieved Dice scores of 0.8888 for WT, 0.7973 for TC, and 0.7331 for ET. Compared with the controlled baseline, QCR therefore produced absolute improvements of 0.0533 for WT, 0.0520 for TC, and 0.0256 for ET.

The results show that the QCR framework improves segmentation performance over the controlled baseline across all three tumor regions. The improvement is also observed for the challenging ET region, where the Dice score increased from 0.7075 to 0.7331.

The fixed-region ablation produced slightly higher WT and TC Dice values than QCR, while both approaches obtained the same ET Dice of 0.7331. The statistical comparison between fixed-region learning and QCR is presented separately in Section 5.5. The Dice performance of the three experimental configurations across WT, TC, and ET is illustrated in Figure 2.

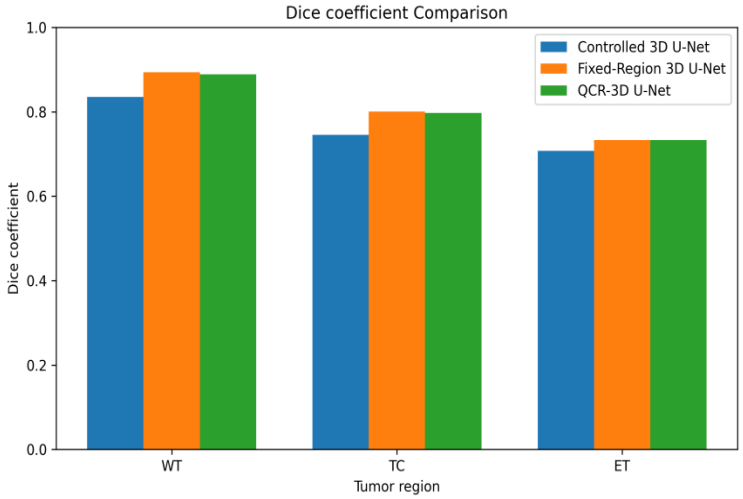


Fig. 2. Dice coefficient comparison across the Controlled 3D U-Net, Fixed-Region 3D U-Net, and QCR-3D U-Net for WT, TC, and ET.

5.2 Detailed QCR Test Performance

The detailed performance of the proposed QCR model was evaluated independently for WT, TC, and ET on the 56-patient held-out test set. The results are reported as mean ± standard deviation across the test subjects.

The complete patient-level performance of the QCR-3D U-Net across WT, TC, and ET is presented in Table 5.

Table 5. Patient-level test performance of QCR-3D U-Net

Region	Dice	IoU	Precision	Sensitivity	Specificity
WT	0.8888 ± 0.1067	0.8120 ± 0.1310	0.8743 ± 0.1124	0.9131 ± 0.1267	0.9970 ± 0.0024
TC	0.7973 ± 0.2500	0.7155 ± 0.2605	0.8183 ± 0.2417	0.8448 ± 0.2111	0.9981 ± 0.0038
ET	0.7331 ± 0.2360	0.6222 ± 0.2427	0.7608 ± 0.2259	0.8084 ± 0.2216	0.9992 ± 0.0009

The QCR model achieved its highest Dice and IoU values for the WT region, with Dice of 0.8888 ± 0.1067 and IoU of 0.8120 ± 0.1310. The TC region achieved a Dice of 0.7973 ± 0.2500, while the ET region achieved 0.7331 ± 0.2360.

For ET, the model obtained a precision of 0.7608 ± 0.2259 and sensitivity of 0.8084 ± 0.2216, indicating that the model detected a relatively high proportion of the enhancing-tumor voxels while maintaining a moderate level of precision.

Specificity remained high across all three regions, ranging from 0.9970 to 0.9992. The larger standard deviations observed for TC and ET indicate greater patient-level variability in segmentation performance for these smaller or more heterogeneous tumor regions.The IoU performance across the three experimental configurations for WT, TC, and ET is illustrated in Figure 3.

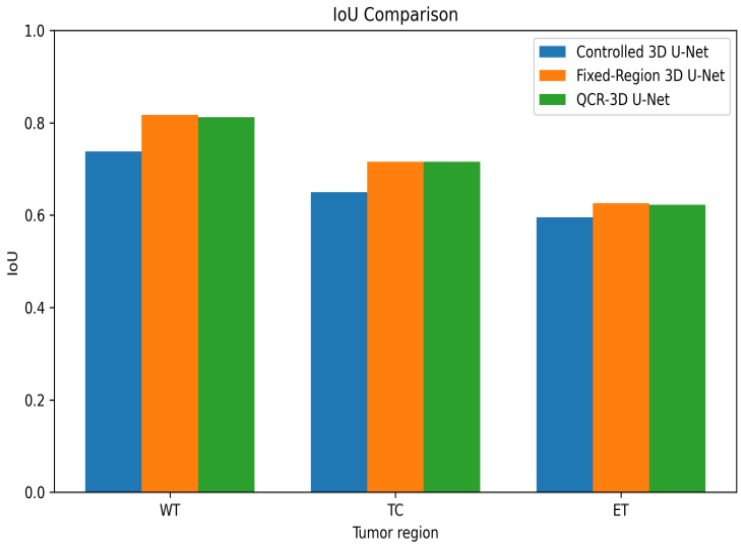
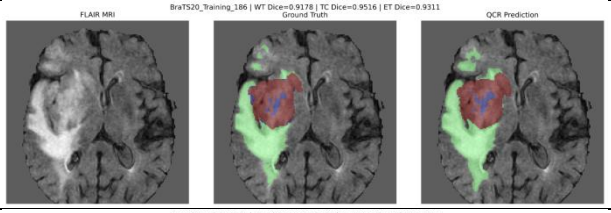


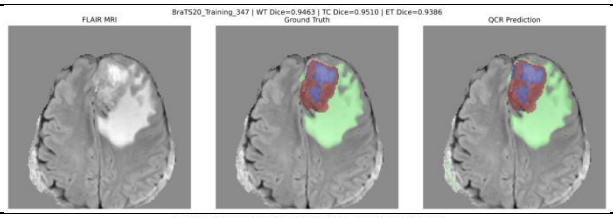
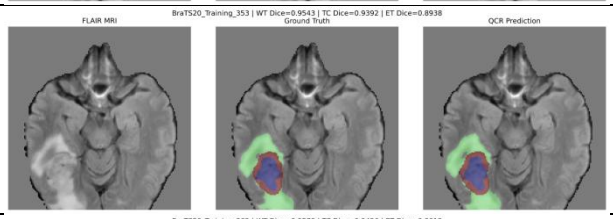
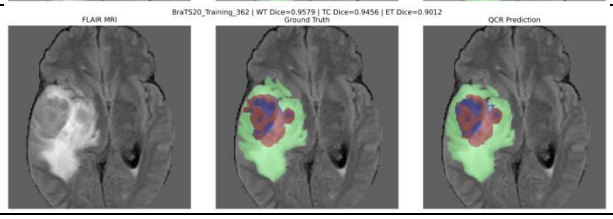
Fig. 3. IoU comparison across the three experimental configurations for WT, TC, and ET.

5.2.1 Patient-Level Qualitative Segmentation Results

To further examine the patient-level segmentation behaviour of the proposed QCR-3D U-Net, representative qualitative results are presented in Table 6. The table provides the patient-wise Dice scores for the whole tumor (WT), tumor core (TC), and enhancing tumor (ET), together with visual comparisons of the FLAIR MRI, ground-truth segmentation, and QCR prediction. Representative patient-level qualitative segmentation results are presented in Table 6, showing the FLAIR MRI, ground-truth segmentation, and QCR prediction for five selected test patients, together with their corresponding WT, TC, and ET Dice scores.

Table 6. Representative patient-level qualitative segmentation results of the proposed QCR-3D U-Net.

S. No.	Patient ID	WT Dice	TC Dice	ET Dice	Qualitative Result
1	Patient_112	0.9491	0.9678	0.9369	
2	Patient_186	0.9178	0.9516	0.9311	
3	Patient_339	0.9253	0.9553	0.9292	

S. No.	Patient ID	WT Dice	TC Dice	ET Dice	Qualitative Result
4	Patient_347	0.9463	0.9510	0.9386	
5	Patient_353	0.9543	0.9392	0.8938	
6	Patient_362	0.9579	0.9456	0.9012	

5.3 Improvement over the Controlled Baseline

The comparison with the controlled baseline demonstrates the effect of introducing the proposed QCR-based learning strategy while retaining the same basic 3D U-Net architecture. The absolute Dice improvements achieved by the QCR model over the controlled 3D U-Net baseline for WT, TC, and ET are summarized in Table 7.

Table 7. Improvement over Controlled Baseline.

Region	Baseline Dice	QCR Dice	Absolute Improvement
WT	0.8355	0.8888	+0.0533
TC	0.7453	0.7973	+0.0520
ET	0.7075	0.7331	+0.0256

The QCR model improved the Dice coefficient for all three tumor regions compared with the controlled baseline. The largest absolute improvement was observed for WT (+0.0533), followed by TC (+0.0520). The ET region improved by +0.0256, increasing from 0.7075 to 0.7331.

These results indicate that incorporating quality-conditioned region-specific learning was associated with improved segmentation performance relative to the controlled 3D U-Net under the same experimental protocol. The improvement in ET is particularly relevant because ET was the lowest-performing region in the controlled baseline.

However, these absolute improvements describe the difference from the controlled baseline and do not by themselves establish statistical significance. The comparison between QCR and the fixed-region ablation is therefore examined separately using paired statistical testing in Section 5.5.

5.4 Fixed-Region Ablation Analysis

The fixed-region experiment was conducted to examine the contribution of region-specific learning independently of the quality-conditioning mechanism. The same controlled 3D U-Net architecture and

experimental protocol were used, with fixed regional weighting applied instead of patient-specific QCR conditioning.

The performance of the fixed-region ablation across the three tumor regions is presented in Table 8.

Table 8. Fixed-Region Ablation Performance on the 56-Patient Held-Out Test Set

Tumor Region	Dice	IoU
WT	0.8933 ± 0.0919	0.8174 ± 0.1261
TC	0.8002 ± 0.2385	0.7160 ± 0.2551
ET	0.7331 ± 0.2457	0.6256 ± 0.2505

Compared with the controlled 3D U-Net baseline, the fixed-region model increased the Dice coefficient from 0.8355 to 0.8933 for WT, from 0.7453 to 0.8002 for TC, and from 0.7075 to 0.7331 for ET. The corresponding absolute improvements were +0.0578, +0.0549, and +0.0256, respectively.

The results demonstrate that introducing region-specific learning produced substantial improvements over the controlled baseline across all three tumor regions. The improvement was also observed for ET, where the Dice coefficient increased from 0.7075 to 0.7331.

The fixed-region results provide an important ablation reference for interpreting the additional contribution of the proposed quality-conditioning mechanism. The direct fixed-region versus QCR comparison is presented in Section 5.5.

5.5 QCR versus Fixed-Region Analysis

To determine whether the quality-conditioned region-specific learning provided a statistically significant difference compared with fixed regional weighting, a paired Wilcoxon signed-rank test was performed using the Dice scores obtained from the same 56 test subjects. The paired comparison between the QCR and fixed-region models, including the mean Dice differences and Wilcoxon signed-rank test results, is presented in Table 9.

Table 9. Paired Wilcoxon Signed-Rank Test Comparing QCR and Fixed-Region Models

Tumor Region	QCR Dice	Fixed-Region Dice	QCR − Fixed	Wilcoxon p-value
WT	0.8888	0.8933	−0.0045	0.1511
TC	0.7973	0.8002	−0.0029	0.7941
ET	0.7331	0.7331	≈ −0.0001	0.7208

The Wilcoxon test showed no statistically significant difference between the QCR and fixed-region models for WT ($p = 0.1511$), TC ($p = 0.7941$), or ET ($p = 0.7208$).

Although the fixed-region model produced slightly higher mean Dice values for WT and TC, the differences were small and were not statistically significant. For ET, both approaches achieved essentially the same mean Dice performance (0.7331).

These findings indicate that the principal performance improvement over the controlled baseline is associated with region-specific learning. The QCR mechanism provides a patient-specific quality-conditioned extension of this regional learning strategy, while maintaining comparable Dice performance to fixed regional weighting.

Therefore, the results do not support a claim that QCR statistically outperforms the fixed-region model. Instead, they demonstrate that quality conditioning can be incorporated into region-specific learning without a statistically significant change in Dice performance under the present experimental setting.

5.6 Precision–Sensitivity Analysis

To further examine the behaviour of the QCR model, precision and sensitivity were compared with those of the fixed-region model using the same 56-patient test set. This analysis provides additional insight into how quality-conditioned regional learning affects the balance between false-positive and false-negative predictions. The precision and sensitivity values of the fixed-region and QCR models across WT, TC, and ET are summarized in Table 10.

Table 10. Precision and Sensitivity Comparison Between Fixed-Region and QCR Models

Tumor Region	Fixed-Region Precision	QCR Precision	Fixed-Region Sensitivity	QCR Sensitivity
WT	0.9003	0.8743	0.8970	0.9131
TC	0.8745	0.8183	0.7886	0.8448
ET	0.8076	0.7608	0.7657	0.8084

The QCR model achieved higher sensitivity than the fixed-region model for all three tumor regions. The increase was observed for WT (0.8970 to 0.9131), TC (0.7886 to 0.8448), and ET (0.7657 to 0.8084). At the same time, QCR showed lower precision for WT, TC, and ET.

For the clinically challenging ET region, QCR increased sensitivity from 0.7657 to 0.8084, while precision decreased from 0.8076 to 0.7608. This indicates that the QCR model identified a larger proportion of true enhancing-tumor voxels, accompanied by an increase in false-positive predictions.

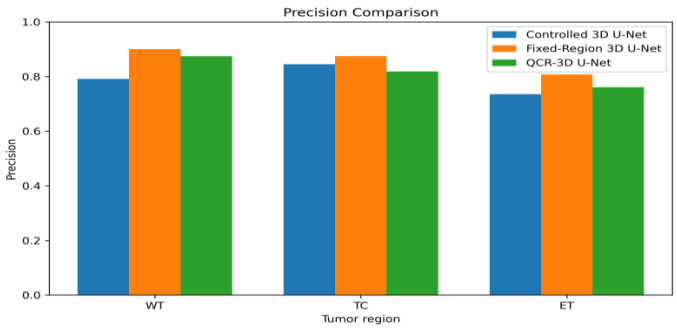


Fig. 4. Precision comparison between the Fixed-Region 3D U-Net and QCR-3D U-Net for WT, TC, and ET.

Overall, these results indicate a precision–sensitivity trade-off associated with quality-conditioned regional learning. The observed behaviour is consistent across WT, TC, and ET and provides an additional interpretation of the comparable Dice performance observed between QCR and fixed-region learning. The precision values of the fixed-region and QCR models across WT, TC, and ET are compared in Figure 4.

The corresponding sensitivity values are compared in Figure 5, showing the observed sensitivity–precision trade-off between the two configurations.

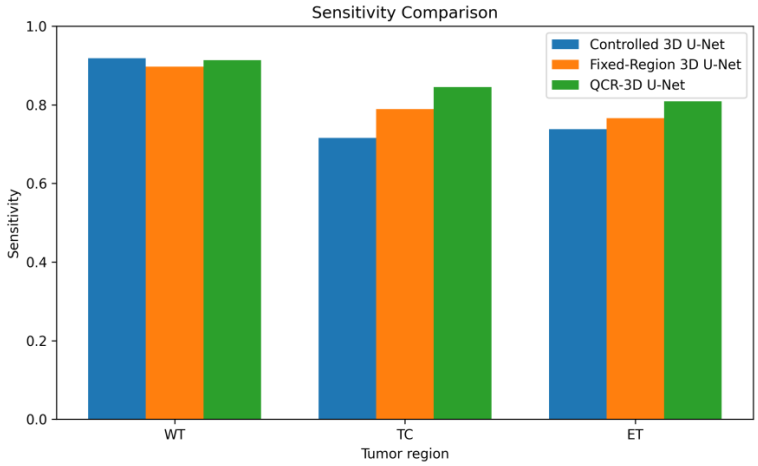


Fig. 5. Sensitivity comparison between the Fixed-Region 3D U-Net and QCR-3D U-Net for WT, TC, and ET.

5.7 Summary of Experimental Findings

The experimental results demonstrate the contribution of region-specific learning and quality conditioning to multimodal 3D brain tumor segmentation. Compared with the controlled 3D U-Net, both the fixed-region and QCR models produced higher Dice scores across WT, TC, and ET regions.

The QCR model achieved Dice scores of 0.8888 ± 0.1067 for WT, 0.7973 ± 0.2500 for TC, and 0.7331 ± 0.2360 for ET, corresponding to absolute improvements of 0.0533, 0.0520, and 0.0256, respectively, over the controlled baseline.

The fixed-region ablation demonstrated that region-specific learning accounts for a substantial part of this improvement. Direct comparison between fixed-region learning and QCR using paired Wilcoxon tests showed no statistically significant Dice differences for WT, TC, or ET. Thus, QCR provides a quality-conditioned extension while maintaining comparable segmentation performance.

The precision–sensitivity analysis further showed that QCR consistently produced higher sensitivity and lower precision than fixed-region learning, including for the challenging ET region. Overall, the findings support the use of quantitative MRI-quality information as an additional conditioning signal within a region-specific segmentation framework, while indicating that further investigation is required to determine whether quality conditioning can produce statistically significant improvements under broader datasets and experimental settings.

6. Discussion

The experimental findings demonstrate that incorporating region-specific learning into a controlled 3D U-Net substantially improves brain tumor segmentation performance, while the proposed Quality-Conditioned Region-Specific (QCR) framework provides a quality-aware extension of this strategy. The results should therefore be interpreted in terms of both the contribution of region-specific learning and the additional role of quantitative MRI-quality information.

6.1 Effect of Region-Specific Learning

The experimental results demonstrate that region-specific learning contributed substantially to the improvement in segmentation performance. Compared with the controlled 3D U-Net, the fixed-region model increased the mean Dice score from 0.8355 to 0.8933 for WT, from 0.7453 to 0.8002 for TC, and from 0.7075 to 0.7331 for ET.

The improvement was observed across all three tumor regions, with the largest absolute gains for WT and TC. The ET region also improved from 0.7075 to 0.7331, indicating that region-specific optimization can provide a measurable benefit for the most challenging tumor component.

These findings suggest that treating WT, TC, and ET as region-specific segmentation targets rather than relying only on a common segmentation objective can better account for their different characteristics and segmentation difficulties. The fixed-region ablation therefore provides experimental evidence that the regional learning component is an important contributor to the overall performance of the proposed framework.

However, the direct comparison between fixed-region learning and QCR showed no statistically significant Dice difference for WT, TC, or ET. Therefore, the results support the contribution of region-specific learning as the primary source of improvement over the controlled baseline, while QCR can be interpreted as a quality-conditioned extension of this regional learning strategy rather than as a statistically superior alternative.

6.2 Contribution of Quality Conditioning

The QCR framework incorporates quantitative MRI-quality information as an additional conditioning signal for region-specific segmentation. Instead of relying solely on the image-derived features of the 3D U-Net, the model uses a 16-dimensional quality representation consisting of CVP, SNR3, CJV, and SNR4 measurements across the four MRI modalities.

The QCR model achieved Dice scores of 0.8888 for WT, 0.7973 for TC, and 0.7331 for ET, compared with 0.8355, 0.7453, and 0.7075, respectively, for the controlled 3D U-Net. Thus, the quality-conditioned framework maintained the improvement obtained through region-specific learning while incorporating patient-specific MRI-quality information into the learning process.

The comparison with fixed-region learning provides an important interpretation of this contribution. QCR and fixed-region learning produced comparable Dice performance, with no statistically significant differences for WT, TC, or ET. Therefore, the present results do not establish that quality conditioning independently produces a statistically significant Dice improvement beyond fixed regional weighting.

Nevertheless, quality conditioning provides a mechanism through which MRI-quality characteristics can influence the regional learning process on a subject-specific basis. This represents a different formulation from using fixed regional weights for all subjects and provides a foundation for further investigation of quality-aware segmentation under larger and more diverse datasets.

6.3 Statistical Comparison with Fixed-Region Learning

A paired Wilcoxon signed-rank test was performed to compare the QCR and fixed-region models using Dice scores from the same 56 test subjects. This analysis was conducted separately for WT, TC, and ET to determine whether the observed differences in segmentation performance were statistically significant.

The obtained p-values were 0.1511 for WT, 0.7941 for TC, and 0.7208 for ET. All values were greater than the conventional significance level of 0.05. Therefore, no statistically significant difference in Dice performance was observed between QCR and fixed-region learning for any of the three tumor regions.

The mean Dice scores were also closely matched. The fixed-region model achieved slightly higher values for WT and TC, whereas both models achieved essentially the same mean Dice for ET. These differences should therefore be interpreted as performance variations rather than evidence of statistical superiority of either approach.

The statistical analysis supports an important interpretation of the proposed framework: the major performance gain over the controlled baseline is associated with region-specific learning, while the QCR mechanism introduces quantitative MRI-quality conditioning without producing a statistically significant change in Dice performance relative to fixed regional weighting.

Thus, within the present 56-patient test cohort, QCR should be interpreted as a quality-conditioned formulation of region-specific learning, rather than as a statistically superior replacement for fixed-region learning.

6.4 Precision–Sensitivity Behaviour

The precision–sensitivity analysis provides additional insight into the behaviour of the QCR model relative to fixed-region learning. Across WT, TC, and ET, QCR achieved higher sensitivity but lower precision than the fixed-region model.

For WT, sensitivity increased from 0.8970 to 0.9131, while precision decreased from 0.9003 to 0.8743. For TC, sensitivity increased from 0.7886 to 0.8448, accompanied by a decrease in precision from 0.8745 to 0.8183. The same pattern was observed for ET, where sensitivity increased from 0.7657 to 0.8084, while precision decreased from 0.8076 to 0.7608.

The ET results are particularly relevant because ET represents the most challenging tumor region in the present experiments. The increase in ET sensitivity indicates that the QCR model identified a larger proportion of true enhancing-tumor voxels, although this was accompanied by reduced precision. This behaviour is consistent with a shift toward detecting more positive tumor voxels at the cost of additional false-positive predictions.

Therefore, the precision–sensitivity results indicate a consistent trade-off introduced by the QCR formulation. This trade-off provides complementary information to the Dice-based comparison and should be considered when interpreting the behaviour of quality-conditioned region-specific learning.

6.5 Significance of Statistical Quality-Feature Selection

The statistical selection of MRI quality features was used to construct a compact and non-redundant quality representation for the QCR framework. Multiple candidate image-quality measures were initially considered, including CV, PSNR, SNR1, SNR2, SNR3, SNR4, CNR, CVP, CJV, and EFC. The selection process considered the variability and distribution of the measures across MRI modalities, redundancy between candidate features, coverage of different quality characteristics, and their relevance to segmentation performance.

Spearman correlation analysis was used to examine redundancy among the candidate quality measures. In addition, a patient-level IQM–Dice relevance analysis was performed using the validation cohort. This provided an additional basis for determining which quality measures were relevant to the segmentation task before the quality representation was finalized.

Based on this analysis, CVP, SNR3, CJV, and SNR4 were selected for each of the four MRI modalities, resulting in a 16-dimensional quality vector. The selected features provide complementary information related to intensity variation, signal characteristics, and image-quality variation while avoiding unnecessary expansion of the conditioning representation.

Importantly, the quality-feature selection was completed before evaluation on the independent 56-patient test set. The selected 16-dimensional representation was then fixed and used as the quality-conditioning input to the QCR model. This separation helps maintain a clear distinction between quality-feature selection using the development data and final segmentation evaluation on the held-out test cohort.

Thus, the statistical feature-selection stage provides a systematic basis for the quality representation used by QCR rather than relying on an arbitrary collection of MRI-quality measures.

6.6 Overall Interpretation

Overall, the experimental findings indicate that the proposed QCR framework provides a practical approach for incorporating quantitative MRI-quality information into region-specific multimodal brain tumor segmentation. The comparison with the controlled 3D U-Net demonstrated improvements in Dice performance for WT, TC, and ET, while the fixed-region ablation showed that region-specific learning accounts for a substantial portion of these improvements.

The direct comparison between QCR and fixed-region learning showed comparable Dice performance, with no statistically significant differences for WT, TC, or ET. Therefore, the present results do not establish an independent statistical performance advantage of quality conditioning over fixed regional weighting. Instead, QCR extends region-specific learning by allowing the regional weighting mechanism to be conditioned on patient-specific MRI-quality information.

The precision–sensitivity analysis further showed that QCR consistently increased sensitivity while reducing precision across the three tumor regions. This behaviour was also observed for ET, where the increase in sensitivity indicates greater detection of true enhancing-tumor voxels but was accompanied by additional false-positive predictions.

Taken together, the results support the feasibility of **quality-conditioned region-specific learning** for multimodal 3D brain tumor segmentation. The findings also indicate that larger and more diverse datasets, together with further patient-level analysis of MRI quality and segmentation behaviour, are needed to determine whether quality conditioning can provide statistically significant improvements beyond fixed regional learning.

7. Conclusion and Future Work

7.1 Conclusion

This study presented a Quality-Conditioned Region-Specific Learning (QCR) framework for multimodal 3D brain tumor segmentation using T1, T1ce, T2, and FLAIR MRI images. The framework integrates a statistically selected 16-dimensional MRI-quality representation with region-specific learning to provide quality-conditioned regional weighting within a 3D U-Net segmentation architecture.

Experiments were conducted on 368 valid BraTS 2020 subjects, using 257 subjects for training, 55 for validation, and 56 independent subjects for testing. The QCR model achieved mean Dice scores of 0.8888 ± 0.1067 for WT, 0.7973 ± 0.2500 for TC, and 0.7331 ± 0.2360 for ET. Compared with the controlled 3D U-Net, these correspond to absolute improvements of 0.0533, 0.0520, and 0.0256, respectively.

The fixed-region ablation demonstrated that region-specific learning was an important contributor to the observed improvement over the controlled baseline. The subsequent paired statistical comparison between fixed-region learning and QCR showed no statistically significant Dice differences for WT, TC, or ET. Therefore, the results do not establish QCR as statistically superior to fixed regional weighting. Instead, QCR provides a quality-conditioned extension of region-specific learning while maintaining comparable segmentation performance.

The precision–sensitivity analysis showed that QCR consistently produced higher sensitivity and lower precision across the three tumor regions. In particular, ET sensitivity increased from 0.7657 to 0.8084, accompanied by a reduction in precision from 0.8076 to 0.7608. This demonstrates the precision–sensitivity trade-off associated with the quality-conditioned formulation.

Overall, the study demonstrates the feasibility of incorporating quantitative MRI-quality information into region-specific multimodal brain tumor segmentation. The findings provide a foundation for further investigation of quality-aware segmentation using larger, more diverse, and independent datasets.

7.2 Limitations

Despite the encouraging results, several limitations of the present study should be acknowledged. First, the experiments were conducted using the BraTS 2020 dataset, and therefore the generalizability of the proposed QCR framework to other datasets, institutions, scanners, and acquisition protocols has not yet been established.

Second, the final evaluation was performed on 56 held-out test subjects from the same dataset. Although the test cohort was kept separate from training and validation, evaluation on independent external or multi-center datasets would provide stronger evidence regarding the robustness of the framework.

Third, the statistical comparison between QCR and fixed-region learning did not show a statistically significant Dice difference for WT, TC, or ET. Consequently, the present experiments do not establish that quality conditioning independently improves Dice performance beyond fixed regional weighting.

Fourth, the QCR model exhibited a precision–sensitivity trade-off, with higher sensitivity but lower precision than fixed-region learning across the evaluated tumor regions. This behaviour indicates that quality-conditioned regional learning may influence the balance between detecting tumor voxels and limiting false-positive predictions.

Finally, the quality representation was constructed from a selected set of quantitative MRI-quality measures and evaluated within the present experimental setting. Further studies are required to determine whether alternative quality measures, larger cohorts, different imaging protocols, or other quality-conditioning mechanisms can provide additional benefits, particularly for the challenging ET region.

7.3 Future Work

Future work will focus on evaluating the QCR framework beyond the current experimental setting and further investigating the relationship between MRI quality and region-specific segmentation. First, the framework can be evaluated on independent and multi-center MRI datasets to assess its generalizability across different scanners, acquisition protocols, and patient populations.

Second, larger datasets can be used to investigate the patient-level relationship between quantitative MRI-quality measures and segmentation performance in greater detail. Such analysis may help identify quality characteristics that are particularly relevant to difficult tumor regions.

Third, future studies can investigate adaptive region-specific learning mechanisms in which the contribution of quality information is learned more explicitly for individual tumor regions, with particular attention to the challenging enhancing tumor (ET) region.

Alternative quality-conditioning strategies, attention mechanisms, and quality-aware feature-fusion approaches can also be explored to determine whether they provide additional benefits while maintaining an appropriate precision–sensitivity balance.

Finally, the segmentation framework can be integrated with quantitative tumor analysis and visualization, including tumor volume estimation, tumor-region characterization, and 2D/3D visualization, to support a more comprehensive computer-aided brain tumor analysis pipeline.

References

1. Z. U. Abidin, R. A. Naqvi, A. Haider, H. S. Kim, D. Jeong, and S. W. Lee, "Recent deep learning-based brain tumor segmentation models using multi-modality magnetic resonance imaging: A prospective survey," *Frontiers in Bioengineering and Biotechnology*, vol. 12, 1392807, 2024.

2. Y. Zhou et al., "Deep learning for multimodal brain tumor segmentation: Architectures, fusion, robust learning, and deployment perspectives," *Computerized Medical Imaging and Graphics*, vol. 134, 102807, 2026.

3. F. J. Dorfner et al., "A review of deep learning for brain tumor analysis in MRI," *npj Precision Oncology*, vol. 9, article 2, 2025.

4. "Advancing Precision: A Comprehensive Review of MRI Segmentation Datasets from BraTS Challenges (2012–2025)," *Scientific Reports*, 2025.

5. W. Aslam et al., "Enhanced brain tumor segmentation in medical imaging using multi-modal multi-scale contextual aggregation and attention fusion," *Scientific Reports*, vol. 15, article 37308, 2025.

6. Z. Li et al., "Unsupervised generative AI for enhancing brain tumor segmentation in multi-center, incomplete real-world data scenarios," *npj Precision Oncology*, vol. 9, article 384, 2025.

7. "A Review of U-Net Based Deep Learning Frameworks for MRI-Based Brain Tumor Segmentation," *Diagnostics*, vol. 16, no. 4, 506, 2026. DOI: 10.3390/diagnostics16040506.

8. R. Preetha, J. P. Priyadarsini, and J. S. Nisha, "Brain tumor segmentation using multi-scale attention U-Net with EfficientNetB4 encoder for enhanced MRI analysis," *Scientific Reports*, vol. 15, 9914, 2025.

9. W. Yang, Z. Li, C. Du *et al.*, "HLNet: high-level attention mechanism U-Net++ for brain tumor segmentation in MRI," *Applied Intelligence*, vol. 55, 692, 2025.

10. Al-Hawasy, J. A. A., & Awad, M. S. (2026). Solvability for quaternary system of elliptic boundary value problems. *Results in Nonlinear Analysis*, *9*(2), 186-201.

11. A. Verma and A. K. Yadav, "Brain tumor segmentation with deep learning: Current approaches and future perspectives," *Journal of Neuroscience Methods*, 110424, 2025. DOI: 10.1016/j.jneumeth.2025.110424.

12. Y. Zhou, J. Kanesan, C.-O. Chow, and C. Xu, "Deep learning for multimodal brain tumor segmentation: Architectures, fusion, robust learning, and deployment perspectives," *Computerized Medical Imaging and Graphics*, vol. 134, 102807, 2026. DOI: 10.1016/j.compmedimag.2026.102807.